

**REPRODUCTION
COPY
IS-4 REPORT SECTION**



*High-Power
Microwave-Tube
Transmitters*

LOS ALAMOS NATIONAL LABORATORY
3 9338 00208 8416

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

This work was supported by the US Department of Defense, Army Strategic Defense Command.

*Edited by Robert Graybill, Technical Communication Services,
for Group IS-1*

Printing coordination by Guadalupe Archuleta, Group IS-9

Cover photo: William North installs a klystron in its modulator housing (LANL Photo No. RN90-03 2008).

An Affirmative-Action/Equal-Opportunity Employer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither The Regents of the University of California, the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by The Regents of the University of California, the United States Government, or any agency thereof. The views and opinions of the authors expressed herein do not necessarily state or reflect those of The Regents of the University of California, the United States Government, or any agency thereof.

*High-Power
Microwave-Tube
Transmitters*

William North

LOS ALAMOS NATL. LAB. LIBS.




3 9338 00208 8416

Los Alamos
NATIONAL LABORATORY

Los Alamos, New Mexico 87545





LANL Photo No. CN60-2171

Technicians at work on a modulator for a megawatt-class dual-klystron RF station.

*This book is dedicated to
Lawrie Eaton,
leader of the RF Technology Group
of the Accelerator Operations and Technology Division of
Los Alamos National Laboratory.
It was his vision that what follows could be of
value to new generations of technologists in this field
(and even the occasional graybeard).
It was at his insistence that the task was undertaken.
If it is of help to anyone, the credit is his.
If it is of no help, the fault is mine.*

Contents


1. Introduction	1
2. The Transmitter	2
2.1 High-power microwave tube	2
2.2 Filament supply	4
2.3 Focus magnet supply	4
2.4 Pulse modulator	5
2.5 Pulsed energy storage	5
2.6 Fault current limiting and charge diversion	6
2.7 Voltage sample	6
2.8 Electronic voltage regulator	6
2.9 High-voltage transformer/rectifier	7
2.10 Step start (or soft start)	7
2.11 Variable voltage device	7
2.12 High-speed contactor or switch-gear	8
2.13 SCR primary controller	8
2.14 Unit substation	8
2.15 Cooling system	10
2.16 Control and monitor	10
3. High-Power Microwave-Tube Transmitter RF Circuit	12
3.1 RF input from the exciter	16
3.2 PIN diode switch	16
3.3 Signal couplers	17
3.4 Ferrite isolator	18
3.5 RF driver amplifier	18
3.6 Another isolator	18
3.7 Another coupler	18
3.8 Power splitter	18
3.9 Waster load	19
3.10 Phase shifter	19
3.11 Time-delay shim	19
3.12 Drive-power-level attenuator(s)	19
3.13 More ferrite isolators	20
3.14 High-power amplifiers	20
3.15 Waveguide arc detector	20
3.16 Forward and reverse couplers	22
3.17 High-power circulator	22
3.18 Output hybrid combiner	22
3.19 Waster power coupler	23
3.20 Waster load	23
3.21 Harmonic filter	23
3.22 Antenna duplexer	24
3.23 Waveguide switch	28
3.24 High-power dummy load	28

4. Transmitter Engineering and the 'Ilities'	29
4.1 Reliability	29
4.2 Maintainability	31
4.3 Availability	31
4.4 Electromagnetic susceptibility	32
4.5 Electromagnetic compatibility	33
4.6 Safety	33
4.7 Built-in test equipment.....	39
4.8 Failure modes and effects analysis	40
4.9 Human factors	40
4.10 Configuration control	41
5. And Now a Word From the Sponsor	43
6. What Makes Resistors Hot?	46
7. High-Voltage Insulation	54
8. Pulse Modulators.....	65
8.1 Cathode-controlled electron gun	65
8.2 Modulating anode	66
8.3 Insulated-focus electrode	67
8.4 Intercepting control grid	67
8.5 Shadow-gridded gun	68
8.6 Bonded-grid gun	68
8.7 E-gun considerations	68
9. Cathode Pulsers: Line-Type Modulators	70
9.1 The line-type pulser	71
9.1.1 The discharge circuit	71
9.1.2 The charging circuit	85
9.1.3 An illustrative example of a line-type modulator	89
9.1.4 The Blumlein discharge circuit	93
9.1.5 The Marx impulse generator	97
9.1.6 Pulse-forming networks	98
9.1.7 Output-pulse transformers	101
9.2 Why do we care about pulse-top flatness?.....	106
9.3 Modulator output pulse-to-pulse variations	118
9.4 Line-type pulser discharge switches	121
9.4.1 The hydrogen thyratron	121
9.4.2 The ignitron	129
9.4.3 The thyristor	131
9.4.4 The reverse-blocking diode thyristor	134
9.4.5 Spark gaps	134
9.5 Practical system applications of half-control switches	136
9.5.1 The hydrogen thyratron	136
9.5.2 The triode thyristor	137
9.5.3 The diode thyristor	142
9.5.4 The ignitron	142
9.5.5 The spark-gap switch	143

10. Cathode Pulsers: Hard-Tube Modulators	146
10.1 The full-control switch tube	146
10.2 Hard-tube modulator topology	148
10.3 Storage capacitors	155
10.4 Vacuum tubes as switch tubes	161
10.4.1 The 7560 triode switch tube	167
10.4.2 The WL-8461 triode switch tube	170
10.4.3 The ML- 6544 triode switch tube	170
10.4.4 The S94000E and 4CPW1000KB tetrode switch tubes	173
10.4.5 The ML-8549 magnetically beamed triode	179
10.4.6 The ML- 8618 magnetically beamed triode	182
10.4.7 The L-5012 and L-5097 Injectron™ beam-switch tubes	184
10.4.8 The 8454H Crossatron™ switch	186
10.5 Load lines	189
10.6 Pulse fall time of hard-tube modulator	191
10.7 Some practical examples of hard-tube modulators	195
10.7.1 The 200-MW triode (WL-8461) hard-tube modulator	195
10.7.2 An 18-MW hard-tube modulator using the L-5097 beam-switch tube ...	198
11. Modulating-Anode Pulsers	200
11.1 Modulating-anode-pulsers topologies	200
11.1.1 Passive pull-up, active pull-down	200
11.1.2 Active pull-up, passive pull-down	201
11.1.3 Active pull-up and pull-down	204
11.1.4 An attempt at simplified amplitude control	205
11.1.5 The grid-catcher type of amplitude control	206
11.1.6 The grid-catcher circuit with potentiometer-type control	207
11.1.7 The gate-catcher type of amplitude control	208
11.1.8 The "quasi-resonant" modulating-anode pulser	210
11.1.9 The modulating-anode power supply/cathode-pulsers hybrid	213
11.2 The self-capacitance of modulator decks	215
11.3 Vacuum tubes appropriate for modulating-anode-pulsers service	218
11.3.1 The 4PR250C/8248 tetrode	218
11.3.2 The 8960 tetrode	218
11.3.2 The Y-847 triode	218
11.3.4 The YU-146 tetrode	219
11.3.5 The TH-5184 tetrode	220
11.3.3 The TH-5188 tetrode	220
11.4 Some representative modulating- anode pulser designs	220
11.4.1 The BMEWS dual-klystron modulator	220
11.4.2 The Millstone Hill radar transmitters	221
11.4.3 Dual-klystron modulating-anode pulser for particle-accelerator RF source	224
11.4.4 The modulating-anode pulsers for the MAR-1 radar system	226
11.4.5 The long-range imaging radar transmitter modulating-anode pulsers ...	229
11.4.5 The modulating-anode pulser for the Haystack Auxiliary Radar (HAX)	232
12. Control-Grid Modulators	234
12.1 Grid-modulator topologies	237

12.1.1	Active pull-up, passive pull-down	239
12.1.2	Active pull-up and pull-down grid pulser	241
12.1.3	Single-switch active pull-up and pull-down grid pulser	241
12.1.4	Cascaded solid-state active pull-up and pull down	243
12.1.5	Hybrid grid pulser	245
12.1.6	"Inductive-kick" single-switch grid pulser	247
12.2	Transmitter applications of grid pulsers	249
13.	Charge Diverters: the Electronic Crowbar	253
13.1	What happens when an electron gun arcs?	253
13.2	Types of crowbar switches	257
13.3	The nature of surge-limiting resistors	261
14.	High-Voltage AC/DC Conversion, or Rectification	263
14.1	Polyphase ac concepts	263
14.2	The three-phase, half-wave rectifier	265
14.3	The three-phase, full-wave (six-pulse) rectifier	267
14.4	The six-phase, full-wave (12-pulse) rectifier	268
14.5	The 12-phase, full-wave, 24-pulse rectifier	271
14.6	Polyphase-rectifier line current	272
14.7	Who cares about harmonic pollution?	277
14.8	Other forms of line pollution	279
15.	Variable-Voltage Devices	282
15.1	The phase-controlled rectifier	282
15.2	The primary SCR controller	284
15.3	The variable autotransformer	288
15.4	The variable-voltage transformer	289
15.5	The step regulator	291
15.6	The induction voltage regulator	292
16.	Electronic Voltage Regulators	295
16.1	Voltage regulators that handle system average current	295
16.2	Voltage regulators that handle system peak current	298
16.3	Voltage regulators that do not handle the full microwave-tube beam current	301
16.4	Regulators that must also be modulators	302
17.	Switchgear and Substations	306
17.1	Switchgear	306
17.2	The unit substation	309
18.	Switch-Mode Electronic Power Conditioning	311
18.1	Switch-mode dc variable-voltage circuits (dc-to-dc converters)	311
18.2	DC-AC inverters	314
18.3	Voltage-multiplier rectifier circuits	322
18.4	High-voltage dc power supply using a quasi-resonant inverter	326
18.5	The electronic high-frequency power conditioner as applied to a microwave tube	329
18.6	An idea too clever not to share (It's not mine, by the way.)	331

19. Thermal Management.....	334
19.1 Heat transfer by radiation	335
19.2 Heat transfer by conduction in solids	336
19.3 Heat transfer by free, or natural, convection in air	338
19.4 Transfer of heat by forced convection of a cooling fluid	339
19.5 Closed-loop water-cooling systems	342
19.6 Vapor-phase cooling, or convection with state change	347
19.7 Multi-phase cooling	351
20. Transmitter Control, Monitoring, and Interlocking	352
20.1 The relay-logic ladder-network control circuit	352
20.2 A solid-state, serial-control, parallel-indication interlock circuit	355
20.3 Integrated-circuit considerations	362
20.4 The programmable-logic controller	365
20.5 High-speed instrumentation for transmitter monitoring	369
20.6 High-speed crowbar-firing circuits	373
20.7 The pulse "balun"	376
Index	380



**HIGH-POWER MICROWAVE-TUBE
TRANSMITTERS**

by
William North

ABSTRACT

High-Power Microwave-Tube Transmitters describes in detail high-power microwave-tube transmitters and the various subsystems that comprise them. Relying on his long experience as a designer of radar systems and rf stations for particle accelerators, the author also imparts lessons he has learned from his work and opinions he has formed on the transmitter design.

1. Introduction

Just as radio and the movies were supposedly killed off by television, the vacuum tube, as we all know, was done in by the transistor. Like the "death" of Mark Twain, however, reports of the death of the vacuum tube have been "greatly exaggerated." Where large amounts of power must be converted from one form to another—especially if the end form is microwaves (that is, frequencies ranging from 1 GHz to 100 GHz approximately)—vacuum tubes still represent viable technology.

The operative word here is "large." What constitutes a large amount of power is related to the square of the wavelength of the end form of converted power. Whereas millions of watts may be required to satisfy such a criterion at the low-frequency end of the microwave spectrum, kilowatts—or even watts!—may be more than adequate to satisfy it at the upper end of the that spectrum.

Despite the earnest efforts of the developers of solid-state devices to make the vacuum tube obsolete, it still rules the realm of high-power microwave. (This is notwithstanding the solid-state designer's "modular" approach to high-power, exemplified by the dictum: "If you can get one watt from a single transistor, you can get a megawatt from a million of them.") High-power vacuum tubes come in many forms, from that hoary workhorse known as the multi-cavity klystron to the most sophisticated wide-band, high-gain, coupled-cavity traveling-wave tube—or even the high-efficiency crossed-field amplifier. These devices still have a role to play in technology.

What follows, then, is dedicated to those technologists who, at some point in their careers, may have to deal with a microwave vacuum tube in order to power a radar system or a particle accelerator, to transmit information or to cure plywood, to kill cancer or even to kill weed seeds before planting. But this book assumes that the reader already possesses a knowledge of microwave tubes themselves. Instead, it focuses on the types of transmitters, or "life-support systems," that use them. To these transmitters, the microwave tubes themselves are the ultimate load.

Although the term "transmitter" will be used throughout this book to identify an overall system whose useful output is high-frequency power, bear in mind that more and more the term is becoming a misnomer; the output of such devices is more frequently being "transmitted" nowhere. Instead, their output is being used internally by an even larger system.

2. The Transmitter

When you put simple things together, you get something complicated.

When you put complicated things together, you get something that probably won't work.

There are those who believe, or at least hope, that their state-of-the-art microwave power tube, which they may have nursed through a costly and time-consuming development process, will come to them complete with a line cord that can be plugged into the wall so that they can be immediately "on the air." Unfortunately for many, this is far from the case.


A typical life-support system for a microwave power tube is called a transmitter and is shown in the block diagram of Fig. 2-1. Not all transmitters require all of the equipment that fills all of the blocks. Each functional block, however, must be evaluated to make sure that every function called out has been satisfied or is not needed. Often, multiple functions can be satisfied by the same subsystem.

2.1 High-power microwave tube

To transmitter designers, a microwave tube is reducible to an electrical load (or loads) to the power-conditioning systems and a thermal load to the cooling system. Once the load requirements have been satisfied and radio-frequency (RF) input has been applied, there is not much that the transmitter designer can do if RF output does not appear.

The microwave tube depicted, a klystron, presents the most rudimentary of all loads to the high-voltage system. Its "body," which comprises the anode and all of the RF internal circuits, is grounded for the well-being of all concerned because the RF input and output connections are made to it. The spent-beam collector can be at body potential, or it can consist of one or many separate stages that are insulated from the body and operated at voltages negative with respect to ground. This is called "depressed-collector" operation. Because the anode is grounded, the cathode must also be operated at a voltage that is negative with respect to ground. As we will see later, this negative cathode voltage may be as great as several hundred thousand volts, either continuously applied or in the form of single-event or time-repetitive pulses. In either case, the klystron, regardless of its size and expense, will behave as a thermionic diode in the space-charge-limited regime, where the current through it will be proportional to the voltage across it, raised to the $3/2$ power (Child's law). The diode portion of the

Figure 2-1. The building blocks of a high-power, pulse-modulated microwave-tube transmitter.



microwave tube is referred to as an electron gun (or, as the French would have it, an electron "cannon"). It determines not only the magnitude of the beam current but the velocity of the individual electrons that make it up. Whereas the current varies as the $3/2$ power of the diode voltage, the electron velocity varies as the $1/2$ power. The total beam power will vary as the $5/2$ power of voltage, and the beam impedance, which is the load impedance the microwave tube presents to the transmitter, varies as the $1/2$ power of voltage.

As we will see later, not all electron guns are simple diodes. Some have insulated anodes called modulating anodes, while others have insulated focus electrodes that are normally near cathode potential, and some have grid structures that are physically and electrically close to the cathodes.

2.2 Filament supply

With the exception of a few exotic cold-cathode devices, electron emission from cathodes is thermionic, requiring that the cathode surface be heated to a temperature in the vicinity of 1000°C . Almost without exception, thermionic cathodes are of the unipotential type, which are indirectly heated (as a slice of bread is indirectly heated in a toaster, rather than by the flow of current directly through it).

The power input requirements of cathode heaters are not in themselves high technology. They usually require less than 1 kW at voltages less than 50 V, ac or dc. What makes the filament supply interesting is that one side of the heater is referenced to the cathode, which, as we remember, may be several hundred thousand volts negative with respect to ground, if for only the duration of a microsecond-long pulse.

2.3 Focus magnet supply

Almost all microwave tubes require a magnetic field of some sort to focus the electron beam in a linear-beam tube or to cause electron spokes to rotate in a crossed-field device. If the tube manufacturer has provided the magnetic circuit as an integral part of the tube—as with periodic-permanent-magnet focusing or even uniform-field permanent-magnet focusing (which is rare but not unheard of)—then this is a box the transmitter designer does not have to worry about, except to keep magnetic material a sufficient distance away from sensitive components. Small linear-beam tubes, such as helix traveling-wave tubes (TWTs), and virtually all crossed-field devices have integral permanent-magnet systems.

Other tubes derive their magnetic fields from solenoids, which are either integral to the tube (wrapped-on solenoid) or separate, where they usually play a role in the physical support of the tube. Energizing such solenoids is the responsibility of the transmitter designer. Single solenoids, such as the ones for millimeter-wave power tubes, can require as much as 10 kW input power. The power must be regulated in current because large variations in solenoid resistance can be expected between its cold and hot states. Some huge multi-cavity klystrons have as many as six or seven separate solenoids arrayed along the length of the electron-beam drift space. Each may require a separate supply with separate adjustment.

2.4 Pulse modulator

In many microwave tube applications there is no need for continuous RF output power. In most radar transmitter and particle-accelerator applications, for instance, RF output is required to be pulsed, so that RF output may last only a few microseconds and recur every few milliseconds. Even the longest pulses last but a small fraction of a second. The duty factor, or duty cycle, of such microwave tube operation may be a small fraction of 1% and is only rarely more than 10%.

Microwave tube power consumption and thermal management considerations dictate that there be some means of turning the electron beam on when RF output is required and turning it off when it isn't. In the case of the power oscillator, the most classic example of which is the magnetron, there is no alternative to turning the beam on and off, because the length of the electron beam pulse is, for all intents and purposes, the length of the RF pulse as well.

Performing the switch function is the role of the pulse modulator. Figure 2-1 shows the most complicated, or at least the most brutish, kind: the hard-tube, full-beam modulator. Vacuum tube *V1*, which can be physically and electrically very large, will have to pass all of the cathode current of the microwave tube during the pulse and block all of its cathode voltage between pulses. Another electronic switch, *V2*, also a vacuum tube, is not commonly used in the application shown but functions as an active pull-up, or "tail-biter," not unlike devices used in integrated-circuit computer-logic elements. Tube *V1* is the active pull-down and *V2* is the active pull-up. The pull-down may have to sink hundreds of amperes, and the pull-up may have to discharge stray capacitance by a hundred thousand volts or more.

Other types of modulators, as we will see, are appropriate for microwave tubes that have additional gun electrodes, such as the modulating-anode and the control-grid tubes.

2.5 Pulsed energy storage

The peak-pulse current in low-duty applications can be hundreds or even thousands of times as great as the average beam current. A practical transmitter, therefore, needs a means of storing energy that is slowly accumulated during the interpulse interval and then rapidly delivering it to the load during each pulse. Storage capacitance is invariably used for this purpose because the current through it can change instantaneously, whereas the voltage across it cannot. To prevent capacitor voltage from changing appreciably during a pulse, it may be necessary to store many times as much energy as is delivered to the load during each pulse (except for a special case, which we will investigate). Energy-storage capacitor banks that store as much as a megajoule of energy (the explosive equivalent of a half-pound of TNT) and several coulombs of charge (the equivalent of commonplace lightning discharges) are by no means uncommon. Many times the energy and charge required for normal single-pulse performance is customarily stored. But this presents a problem. If there should be a microwave tube malfunction in the form of an internal electron-gun arc—and in most cases this is a "when" rather than an "if" scenario—all of the stored energy could be

dissipated in the arc, unless something is done about it.

2.6 Fault current limiting and charge diversion

Energy can only be dissipated in resistance. Fortunately, the resistance of an arc in a vacuum between copper electrodes, which is what we have in a microwave tube electron gun, is small and inversely proportional to the current through it, up to several thousand amperes of arc current. Therefore, external resistance, which is large with respect to the arc resistance but small compared to the ratio of normal pulse voltage to current, will limit arc current and dissipate all but a tiny fraction of the stored energy. Supplying this external resistance is not an option. You do not have a usable transmitter without it!

Often current-limiting resistance is all that is required to prevent tube damage resulting from internal high-voltage arcing. In other cases, where the microwave tube is deemed more fragile, a charge diverter, usually called an electronic crowbar, is used. (It is most often a nervous tube vendor who claims that the company's warranty will be void if your system does not have a crowbar to provide this safety feature. However, if your system does have a crowbar and the tube fails anyway, the self-interested vendor will most likely claim that your crowbar did not work.) As shown in Fig. 2-1, the crowbar is nothing more than a device that shunts the high-voltage system within microseconds of the detection of fault current in the microwave tube. This arc does not share the stored energy, but it does share the stored charge. The crowbar is connected to a tap on the total surge resistance. The resistance between the capacitor bank and the crowbar determines the current through the crowbar arc. The resistance between the crowbar and the microwave tube assures that the crowbar takes most of the current and charge by making the total impedance of the path through the faulted tube much greater than that through the fired crowbar.

2.7 Voltage sample

If for no other reason than to provide the control-and-monitor system a proportional sample of voltage, there will be a need for a voltage divider connected across the high-voltage bus. It may also be the source of an error signal for an electronic voltage regulator, in which case it often will include frequency-compensation components to improve its transient response and ripple suppression.

2.8 Electronic voltage regulator

In all precision transmitters (especially those using TWTs as RF amplifiers), some form of regulation of the dc high voltage will be required. In the type illustrated, the linear (variable resistance) electron tube regulator is the most precise and has the best transient response. It is also the most expensive and inefficient. If a voltage regulator is to be used, there are less expensive and more efficient alternatives, which will be discussed later. As we shall also see, there are situations when voltage regulation is required on the load side of the capacitor bank. In such cases, the regulator must handle the peak-pulse current instead of the time-averaged current while maintaining precise output voltage. The linear

vacuum-tube regulator is mandatory for such service.

2.9 High-voltage transformer/rectifier

This block is the beginning of the high-voltage dc system and the end of the primary-power ac system. This is true even in the so-called dc-to-dc-converter power supplies because they have a self-generated ac link within them. In large power systems, the ac service is usually three-phase. The transformer part of the block steps up the primary ac and often converts it to 6-, 12-, or even 24-phase input to the high-voltage rectifier. A well-designed and executed high-voltage dc loop is referenced to ground at only one point: the body of the microwave vacuum tube. Such a high-voltage loop is like electrical code house wiring, except in reverse. House wiring comprises black (hot), white (neutral), and green (safety ground) conductors. The neutral is grounded only at the service entry and nowhere else throughout the house. The microwave-tube dc loop has the negative cathode-potential lead (black equivalent), which must be insulated for tens or even hundreds of thousands of volts from ground; the positive current-return lead (white equivalent), which is isolated from ground but needs to be insulated for only a few hundreds or thousands (transient) of volts; and the equivalent of the green wire, which is the grounding of the individual metal enclosures that house the transmitter equipment. Remember, ground is only ground when there is no current through it.

2.10 Step start (or soft start)

Given that an electrically large energy-storage capacitor bank is a requirement of a pulse-modulated, high-power system—especially if the duty factor is low (high peak-to-average ratio)—it is necessary to face the problem that the initially uncharged bank is a virtual short circuit for the system that electrically precedes it. Means must be provided to limit inrush of the charging current. The simplest form of such limiting is shown in Fig. 2-1: the step-start, which comprises resistance of high enough value to keep current within system rating, energy-handling capability equal to the eventual amount stored in the bank, and a relay that short-circuits it when the bank is fully charged. As resistance, it can be deployed on either the ac or dc side of the transformer/rectifier. If an electronic voltage regulator is used on the input side of the bank, it can be programmed in conductivity to perform the soft-start function. When deployed as shown on the ac side, the resistance can be supplanted by either series ac capacitance or line inductance to increase the ac system source impedance without power dissipation.

2.11 Variable voltage device

In most, if not all, transmitters, it is anywhere from comforting to essential to be able to vary the effective amplitude of the ac input to the high-voltage transformer/rectifier. This is a function that could be performed by the electronic regulator, except that it effectively varies the dc output instead of the ac input. (Unless it happens to be part of a switch-mode dc-to-dc converter, in which case it will vary the dc input to the internal ac generator before it reaches

the transformer/rectifier.)

As we will see, the variable device can be as simple as the variable auto-transformer, or "Variac" in the common parlance, or as complicated as a polyphase silicon-controlled-rectifier (SCR) primary controller. It can also serve the function of "soft-start," supplanting the "step-start" function, or it can even be a true high-voltage SCR acting directly in the dc link.

2.12 High-speed contactor or switch-gear

The term "high-speed contactor" is actually a misnomer in this application. What is required in this block is a high-speed disconnect. When the microwave tube internally arcs and/or the electronic crowbar fires, everything on the line side of them is short-circuited, just as if a physical crowbar had been thrown across the high-voltage dc. Short-circuit current, limited only by the internal impedances of the system (predominantly leakage inductances of the transformers), will flow until the circuit is broken by the action of the high-speed disconnect. Even then, current will not stop until the first current zero is reached after the line circuit is physically broken, because the contacts will arc across, even in vacuum. If a primary SCR controller is used for the variable-voltage function, it will double as a solid-state relay. The gating of the SCRs can be halted with electronic speed much faster than even the fastest electromechanical device. Even so, conduction will once again continue until the first current zero is reached. In a three-phase system, the longest interval between the earliest termination of gate drive and the next current zero is one-third cycle.

2.13 SCR primary controller

The SCR primary controller, or "light-dimmer," can be programmed to perform the functions of high-speed disconnect, output voltage variability, soft start, and even regulation of capacitor-bank dc voltage. Modern high-voltage dc systems increasingly tend to use them. As will be discussed, the SCR primary controller affects all of these functions through phase-controlled conduction in high-efficiency, true switch-mode fashion.

2.14 Unit substation

A unit substation indeed! No, I do not mean the type of electrical substation that is surrounded by a chain-link fence and danger signs and that might be found on the outskirts of town. But I do mean a miniature version of such an installation that performs the same basic functions: circuit protection and voltage reduction.

The first such function, working our way backwards from the transmitter as shown in Fig. 2-1, is carried out by the circuit breakers. A unit substation is designed to accommodate a breaker panel and is, in fact, the most convenient and logical place for one. The second function is to provide a step-down transformer to reduce the subtransmission-line voltage, which for older installations in the USA was likely to have been 4160 V and now 13,200 V. Common secondary voltages in substations are in the 400-600-V range, especially if a primary SCR controller is used with the high-voltage dc power supply. But note that if

the average volt-ampere demand is in the MVA range, line currents can become so large that fault currents resulting from high-impedance faults can hardly be distinguished from normal full-load line currents. If there is not a sufficient ratio between the currents produced by high-impedance faults and normal load currents, circuit breakers will either not trip at all or will require too great a time interval to do so. The result is often an electrical fire, which will please no one. In such situations, a higher secondary voltage is called for (perhaps as great as the incoming subtransmission voltage) to proportionately lower normal line current and increase the fault/normal ratio. But if the substation transformer has unity turns ratio, you might ask why even bother with it. The answer has to do with the *other* rating of a circuit breaker.

We are all familiar with the primary rating of a circuit breaker. If, for instance, we plug too many hair dryers and popcorn poppers into a typical 15-A household outlet, the breaker is going to trip sooner or later—the tripping time will be inversely proportional to the overload ratio—when the total load current exceeds 15 A. Suppose, however, that instead of plugging in too many appliances we plug in a strip of copper, shorting the outlet and thus producing what is called a “bolted” fault. Now we have to worry about the other rating of the breaker, which is how much fault current it can be expected to break. This characteristic is called its interrupting rating. The household-service 15-A breaker may have an interrupting rating of 10,000 A. If the fault current is greater than that, the breaker will surely trip—sooner, rather than later—but it will *not* open the circuit. What assurance do we have that the impedance of the fault will limit the fault current to less than the breaker’s interrupting rating? The answer is none. What *will* limit the fault current is the *source* impedance. But how do we know what the source impedance will be? We have now arrived at the least-understood, but perhaps most-important, function of our unit substation.

The substation’s power transformer, if it is properly sized, will be matched to the maximum volt-ampere demand of the system (plus a modest safety factor added in for comfort). The internal impedance of this transformer consists, for the most part, of the leakage inductance between primary and secondary windings. Leakage inductance, then, is a measure of how much of the flux produced by the primary winding fails to link with the secondary winding. The value of the line-frequency reactance of this leakage inductance is between 5% and 10% of the load impedance at full rating and is typically greater than the resistive component of transformer impedance (so the resistance can be ignored). So secondary fault current, even into a zero-impedance fault, will be between 10 and 20 times the normal full-load current. Often the equivalent of a unit substation is already provided to you as part of the facilities of the physical plant in which the transmitter is installed. If so, the most important feature of the primary power provided, other than its adequacy to provide enough volt-amperes to run the transmitter, is its internal impedance, normalized to the load base. It is up to the designer to find out its internal impedance and to do something about it if there is a problem.

Finally, at the front end of the substation it is comforting to have a fused, load-break manual disconnect switch, which allows a positive, no-nonsense dis-

connect from the incoming primary power. Such a disconnect is an essential feature for transmitter maintenance.

2.15 Cooling system

It is by no means unheard of that the difference between a successful and unsuccessful transmitter design is the cooling system. Often too little attention is paid to the mundane act of cooling electronic components because it is not "high-tech," or at least does not appear to be. Unless they are 100% efficient, all active systems, electrical or mechanical, convert a portion of their input energy to heat. This heat will be dissipated, one way or another. The purpose of a well-designed cooling system is to remove this heat while keeping the temperature rise of the heat source within acceptable limits.

There are some obvious temperature benchmarks. Copper, for instance, melts at 1073°C. Therefore, it is a good idea to keep the inner surface of the microwave tube beam collector below this temperature. Silicon junctions have severely shortened life expectancies above 300°C. (In fact, MIL-HDBK-217, the reliability bible for military systems, does not accept a reliability prediction using its guidelines if the junction temperatures exceed 175°C.) All chemical reactions that govern the life span of many, if not all, electronic components are governed by the Arrhenius equation, which states that the speed of the reaction is a log-normal function of the absolute temperature. The lifetime of silicon junctions, for instance, is halved for each 10°C increase in their temperature. The life expectancy of thermionic emitters in vacuum tubes is also a function of operating temperature. (But that problem is outside the function of a transmitter's cooling system because hot thermionic emitters are a necessary evil.)

Later, quantitative relationships in cooling systems will be discussed, but for now the cooling system is depicted as a closed-loop liquid system with a liquid-to-air heat exchanger. This is because water is the most effective coolant. Some high-power microwave tube manufacturers state their devices are designed to be cooled by conducting away their heat through ambient air. In these cases, be aware that what the tube is mounted to is expected to have the properties of an infinite heat sink (that is, it can conduct away an infinite amount of heat without itself rising in temperature).

2.16 Control and monitor

At last we have reached the nervous system of the transmitter, whose nerve endings are in the form of an array of sensors that sample conditions throughout the transmitter. These sensors include voltage dividers; current monitors; optical and acoustic sensors; temperature, pressure, and flow sensors; and even closed-circuit TV cameras. The control function, or brain, of the transmitter includes the operator, although modern high-power RF sources are less and less dependent on human intervention for their control and rely increasingly on computer-based industrial programmable logic control (PLC) systems. Such systems perform ladder-logic and provide, if you will, the hand-eye coordination needed to successfully operate the transmitter. They also create the computer-screen graphics that make up the operator-machine interface. These graphics are far more infor-

mative and interpretable than the pilot lights and analog meters that are still, unfortunately, more standard than not in the industry.

Yet any control system must have the electronic equivalent of a conditioned-reflex, or knee-jerk, reaction that can in emergencies bypass reasoned thinking or programmed logic in order to do the right thing, every time, as quickly as possible. There are situations where such quick reactions are required. One obvious example that we have discussed already is the firing of the electronic crowbar in response to the detection of a high-voltage arc in our high-power microwave tube. This must occur within a few microseconds. Another is the automatic interruption of the primary power to the high-voltage dc power supply. These circuits, then, must be hard-wired so that their logic is infallible. The PLC, in its own good time, may take note of these involuntary reactions, update the status graphics accordingly, and perform whatever subsidiary activities it has been programmed to do, including automatic reset and restart.

3. High-Power Microwave-Tube Transmitter RF Circuit

Never do anything at high level that you can do at low.

Figure 3-1 shows what might be considered a typical transmitter for a high-power microwave-tube radar system that uses a pair of microwave tubes whose outputs are combined in a single-output waveguide channel. All these transmitter components, if multiplied, would be appropriate even for a phased-array transmitter, except the outputs of the high-power amplifiers would not be combined in a waveguide but would be directed to separate antenna elements, and the phase control would be variable with electronic speed. The interconnections shown are also expandable in order to accommodate many more parallel output devices, which, as we will see quite dramatically later, do not have to grow as powers of two.

The system shown has been implemented in rectangular (1:2 aspect ratio) dominant-mode ($TE_{0,1}$) waveguide. At lower microwave frequencies the components and interconnections are often made of coaxial transmission lines, which can grow in diameter and power-handling capacity until the mean circumference between conductors is equal to a wavelength at the highest operating frequency where the first waveguide mode that is *not* the dominant TEM mode is possible. (This mode is the transverse electric wave, $TE_{1,1}$.) Inner-conductor heating limits coaxial-line average-power capability, which can be enhanced by cooling (but not very easily). Voltage breakdown and/or tracking along the dielectric supports between inner and outer conductors usually limits peak-power handling, which can be enhanced by pressurization or use of a dielectric gas, such as nitrous oxide or sulfur hexafluoride.

Almost all high-power systems that operate above approximately 400 MHz use dominant-mode rectangular (1:2 aspect ratio) waveguide transmission lines and RF components. As average-power levels increase, waveguides, even though they are designed for low loss, can dissipate enough heat to require cooling. Natural convection air cooling can be enhanced by affixing metal fins transverse to the direction of the waveguide run. (Such fins also stiffen the broad wall, minimizing its deflection if waveguide pressurization is needed.) Eventually, as average power is increased, water cooling of the guide is required. This is usually accomplished by soldering or brazing copper tubing to the copper waveguide along the broad wall, where the RF currents are the highest.

RF breakdown in waveguides operated at high power is quite common, even when the electric-field intensity produced by the peak power is but a fraction of the theoretical breakdown value. The culprit in such cases is almost always imperfect mating of the flanges that connect waveguide sections. If the metallic contact around the entire periphery of a waveguide joint is not complete, the RF




Figure 3-1. High-power microwave-tube transmitter RF circuit.

current distribution will not alter itself to avoid the spots with questionable continuity but will rather ionize the air within the void, producing great local heating and a supply of free ions. This phenomenon is called a series arc, which does not stay just a series arc for long, however. A cross-guide arc, which is a very troublesome kind, almost always immediately follows. The corrective solution for cross-guide arcs is not complicated: make sure there are no peripheral voids. Tighten up the bolts and, as part of the original design, make sure that there are enough bolts in each flange to make void-free tightening possible in the first place. (Even better is to minimize the number of flanges.) Where a number of components are closely coupled together, it can pay long-term dividends to fabricate them as a group with no intermediate flange connections. (If the transmitter is to be manufactured in high volume, there may be no cost penalty associated with such a design strategy either.)

As the operating frequency of the transmitter goes up, so does waveguide loss rate. For instance, in the W-band, which encompasses the 94-to-96-GHz moisture-absorption dip band, the loss rate of the dominant-mode guide, WR-10, is approximately 1 dB/ft. At this loss rate, the antenna feed does not have to be very far from the transmitter output before you lose the first 10 dB of system sensitivity. This loss applies not only to the transmitter power before it reaches the antenna feed but also to the received power before it reaches the low-noise preamplifier. In such situations, over-moded sections of special waveguide are used. Often these are circular waveguides, which can have low-loss properties if the run is perfectly straight.

The ultimate low-loss, high-power waveguide is actually no waveguide at all—or at least no waveguide with walls. It is called a beam waveguide and has been used for over 20 years in C-band satellite-communication earth stations. These stations use the beam waveguide to direct power from a transmitter on the ground to an antenna feed on a tower. A beam waveguide uses quasi-optical techniques. From the transmitter output, a Gaussian beam is launched from a “scalar” horn, which has circular cross-section and a corrugated inner surface. The beam that emerges has circular symmetry and a Gaussian distribution for any cross section. The beam is not of constant diameter, and after reiteration and redirection by what is called a “clamshell pair” of reflecting mirrors, it converges toward a “beam waist” and then diverges again, as illustrated in Fig. 3-2. Only at the beam waist positions are the phase centers of all of the rays in a plane perpendicular to the center ray. The parameters of beam width, effective radius of the phase centers, and distance from beam waist are illustrated in Fig. 3-3. Beam width, $\omega_{(z)}$, is defined as

$$\omega_{(z)} = \omega_0 \sqrt{1 + \left(\frac{z}{\delta_0}\right)^2}$$

and the phase center, $R_{(z)}$, is defined as

$$R_{(z)} = z + \frac{\delta_0^2}{z},$$




Figure 3-2. Basic properties of beam-waveguide RF transmission system.

where

$$\delta_0 = \frac{\pi\omega_0^2}{\lambda}.$$

Note that more than one mode is possible.

Only since 1990 has beam-waveguide technology been expanded to serve all of the functions of a sophisticated radar-tracking feed system. This feat was first achieved by the ALCOR/Millimeter-Wave Radar, which operated at a center frequency of 35 GHz. Figure 3-4 shows a simplified three-dimensional schematic diagram of the quasi-optical microwave components developed by MIT/Lincoln Laboratory for the project. Note that the polarization filters are grids of wire stretched across the beam wave space. They are designed to reflect one sense of polarization and be transparent to the orthogonal sense. The circular polarizer is more like a wide-open Venetian blind whose width creates the length of waveguide that generates a quadrature component from a linearly polarized wave. The 45° Faraday rotator performs the function of an antenna duplexer for the oppositely polarized received signal, which, without the total of 90° round-trip phase shift, would end up in the transmitter output waveguide.




Figure 3-3. How width of beam and phase-center radius vary with distance from beam waist.

The rest of this chapter will briefly describe the transmitter components identified in Fig. 3-1.

3.1 RF input from the exciter

The RF input from the exciter is the transmitter's interface with the rest of the complex system, which can be a radar system, a communications system, or the RF source for a particle accelerator. When properly defined, the RF input will tell you what signal level to expect and whether the signal is broadband or has complex modulation. It will also define the signal's spectral and temporal characteristics.

Typical input interface signal levels are between 100 mW and 1 W (-10 dBW and 0 dBW). The people supplying you with this signal (often called "exciter people") will want to know from you—or will specify to you—the transmitter's complex load impedance. If they are clever, experienced, or both, they will insist that the sign of that impedance always be positive. (As will be explained later, it is by no means impossible for a transmitter to reflect more power than is incident upon it and yet operate quite normally, depending upon the *source* impedance of the exciter.)

3.2 PIN diode switch

This is a high-speed shut-off switch capable of interrupting, or greatly attenuating, the RF input signal within fractions of a microsecond. Its purpose is not to produce pulse modulation of the input RF signal—that is in the domain of the exciter—but to prevent damage to the microwave tube or tubes in the event of

high reflected output power or waveguide arcing. The PIN diode switch is a necessary, but not always satisfactory, reactor to a fault condition. If, for instance, the output anomaly caused the microwave tube to self-oscillate, then turning off the RF drive, no matter how rapidly, will do no good. In this situation it is necessary to kill the microwave tube beam current as well, either by terminating a pulse modulator output pulse or by firing an electronic crowbar to dump the entire cathode voltage.

3.3 Signal couplers

Because signal couplers can be designed to house RF sensors, transmitter designers often say that you can't have too many signal couplers. Sensors are usually capacitive or inductive. The simplest of capacitive (paddle) or inductive (loop) probes, which barely protrude into the wave space, can often provide all of

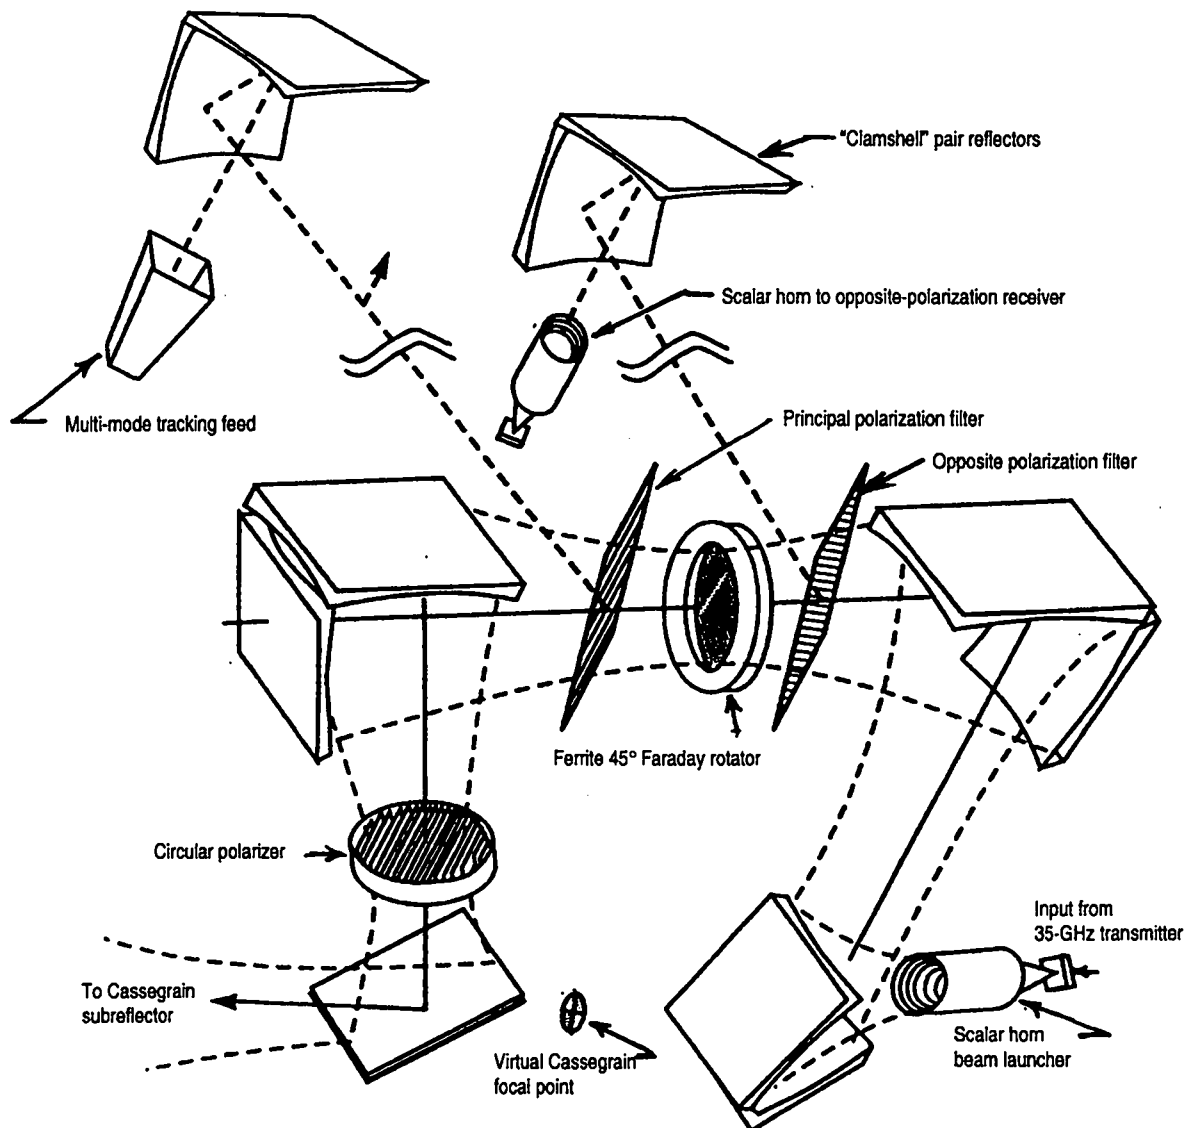


Figure 3-4. Simplified schematic of beam-waveguide components of tracking radar feed.

the information a transmitter designer needs: has the signal gotten this far or not? Often the awareness of the mere presence of RF field activity is sufficient for the designer. There is a special type of signal coupler known as the directional coupler, which can differentiate between forward and reflected RF vectors. There are instances where this kind of information can be useful to the designer. For instance, the first signal coupler illustrated in Fig. 3-1 may very well have to be dual-directional because the designer may want to know not only if RF power is getting past the PIN diode switch but how much of it is being reflected.

3.4 Ferrite isolator

This device may be an isolator if reflected power is small enough to be absorbed by the ferrite material of the device, or it may be a junction circulator if the power is too great to be absorbed by the ferrite. Circulators have three ports—input, output and load—with the reflected power being diverted to the load port. Remember that in some special cases the reflected power absorption or diversion capability of this device may actually exceed the incident power. In virtually all cases, however, this device will be responsible for meeting the input impedance requirement.

3.5 RF driver amplifier

Except in rare instances, the RF input from the exciter will be inadequate to drive the high-power amplifier stages directly. A driver or intermediate amplifier will usually be required. This could be a cascade or distributed-transistor amplifier, a reflection-type impact-avalanche-transit-time (IMPATT) diode amplifier (or locked-oscillator amplifier), a helix-type or low-power coupled-cavity TWT, or, in more mature designs, an amplifier comprising triodes and/or tetrodes that have cavity-type RF circuits. Typical drive amplifier power output ranges from 10 to 100 W (+10 dBW to +20 dBW).

3.6 Another isolator

This isolator functions like the isolator of 3.4, except here it isolates the driver amplifier output from high-power amplifier reflections.

3.7 Another coupler

This coupler functions like the coupler of 3.3, except here it is one stage further in the amplifier chain.

3.8 Power splitter

Because we will be driving two high-power amplifiers (HPAs) from a common drive source, we will need to split the drive line into two channels of equal amplitude. Several types of components can do this, but they differ with respect to the phasing of the output signals. It is not necessary that the outputs have identical phasing, so long as a device with similar phase-shift characteristics is used to recombine the channels after high-power amplification. (If the system is narrow-band or fixed-frequency, even that is not necessary.) Often, if not always, the splitter of choice is a short-slot-sidewall, 3-dB, 90°-hybrid coupler. It is reasonably broadband both in power-split ratio and 90°-output-phase difference, and it has high power-handling capability, which, although not important at this

point in the chain, is important if we use a similar device for recombination. There is another, more subtle advantage to the 90° output-signal phase difference. It can be reasonably expected that the high-power amplifiers will have imperfect input matches and will reflect some of the incident drive signal. Even though the amplifiers may have imperfect input matches, if their reflection coefficients are similar—or better yet, identical—the reflected signals from them, when passing through the 90° hybrid in the reverse direction, will combine in the hybrid waster load (3.9) and will cancel at the input port, in which case the isolator (3.6) will have nothing to do.

3.9 Waster load

This is the resistive RF load mentioned above. Its power rating should be equal to the driver amplifier output-power capability in case of total reflection from the HPA inputs, which is by no means impossible.

3.10 Phase shifter

This device, often motor-driven, is used by an operator or an automatic servo system to maximize the combined RF output from the dual-channel system by making the phase delay of the two parallel channels as identical as possible at the operating frequency (or near the band center if it is a broadband system). It is desirable, but less than crucial, that phase changes do not produce changes in attenuation.

3.11 Time-delay shim

If the system is to operate over a significant band of frequencies, the channels must have equal time delay as well as equal phase delay at a single frequency. If the time delays are not equal, phase delay will accumulate more rapidly in one channel than the other as frequency is increased. (Time delay can be expressed as the quotient of incremental phase divided by incremental frequency, or $\Delta\text{phase}/\Delta\text{frequency}$.) This condition, known as "phase runout with frequency," can be statically corrected by replacing the high-power amplifier tubes with lengths of transmission line that approximate their time-delay contributions and by adding time-delay shims or incremental sections of transmission line to the channel that needs them.

3.12 Drive-power-level attenuator(s)

These attenuators are adjusted during initial system operation and only periodically thereafter, usually following the replacement of an HPA or drive-stage tube. These attenuators assure that both HPA tubes are operating at their maximum power outputs, even if their power outputs are not the same. Although an amplitude imbalance will produce some waster-load power, the loss in total combined output power will be greater than the loss caused by amplitude imbalance if the output of the higher-powered tube is deliberately reduced to match that of the lower-powered one.

3.13 More ferrite isolators

Although isolators in this location are quite common, they are clearly redundant with the one at (3.6). At first glance, it would appear that the isolator at (3.6) would have to have twice the reverse-power handling capability as either of the isolators at (3.13). This overlooks the action of the input 90° power splitter, however. The statistical likelihood is that more than half of the total reflected power from the two HPAs will be diverted to the hybrid waster load, in which case the (3.6) isolator sees less reflected power than either of the (3.13) isolators. If the HPA reflection coefficients are equal and the incident drive signals are equal, the (3.6) isolator does not see any reflected power at all, unless the (3.10) phase shifter has been grossly maladjusted.

3.14 High-power amplifiers

These HPAs may be any coherent power amplifiers. The classic examples are the klystron and the TWT, which, at the highest power levels, is usually of the coupled-cavity variety. If we stretch the concept of coherency a bit, we could also use crossed-field amplifiers, which many feel act more like locked oscillators than true amplifiers. (Speaking of locked oscillators, there is no fundamental reason why injection-locked magnetron oscillators could not be accommodated in these applications, except that the required input and output ports would be obtained by using a circulator connected to the single port of the oscillator.) As will be seen later, it is essential that both amplifier tubes be operated from the same high-voltage cathode bus and, if practical, from the same pulse-modulator bus as well.

3.15 Waveguide arc detector

The only important damage that an arc in the waveguide system can cause—apart from interrupting the orderly flow of RF power to its intended destination—is to the output window of the microwave vacuum tube. The most serious form of such damage is cracking, which usually results in a loss of vacuum within the power tube. The most frequently used protection against such arcs involves the use of an optical arc detector. Moderately fast (on the order of a microsecond) photo-sensitive solid-state diodes are the most commonly used transducers. They are positioned “to see” the window through a small hole in the wall of a waveguide bend that is in line with the window. The sensor mount is in the form of a small-diameter tube that acts as a waveguide beyond the cutoff. The tube prevents RF leakage either to the diode, which may be sensitive to it, or beyond. When the waveguide dimensions themselves are too small to permit the use of such a tube, a small-diameter fiber-optic cable is used to couple the light energy into the photodiode.

Perhaps the highest-performance optical arc detector design, however, uses a photomultiplier vacuum tube as the transducer. In its most sophisticated form, the photomultiplier is continuously self-checking. This is how it works. Through another aperture in the waveguide wall a low-intensity light illuminates the surface of the window. The reflected “background” illumination produces continuous photo current in the photomultiplier and a dynode output voltage proportional to it. The logic circuitry that the tube drives has two threshold levels: a

low-level one, which must be exceeded by the continuous self-check output; and a high-level one, which is below the dynode output signal under arc conditions. If the output signal falls below the lower threshold, a fault in the circuit is indicated. If the signal output exceeds the upper threshold, an arc is indicated.

Other means have been used to detect arcs. Acoustic arc detection has been successfully employed in applications requiring low frequency (L-band and below) and long pulse duration (~ 1 ms)—and this includes continuous wave (CW), of course. Acoustic transducers are nothing more than loudspeakers acting as microphones. They are mounted to the waveguide broad walls. The same kind of matching transformer that is used to connect speakers to a "constant-voltage" bus in public address systems is often used to step up the feeble signal present directly at the terminals. Figure 3-5 is the schematic diagram of the protective circuit logic associated with an acoustic arc detector. It comprises a single ML-339 quad-voltage comparator. The circuit has bipolar thresholds to provide response within the first half-cycle, regardless of which direction the speaker cone starts to move in response to the sound caused by the arc.

Neither optical nor acoustic means are easily applied to coaxial transmission line systems, however. Instead, systems for coaxial lines have been employed that sense RF drive power, power amplifier beam current, and forward RF output power. These signals are sent to logic circuitry. If the first two signals are present and the third is not, an output RF fault is inferred and the RF drive is inhibited. Needless to say, in such systems a small delay must be incorporated in the search for RF output before the RF input is shut off. A somewhat similar scheme was once used in a family of high-power CW klystron amplifiers. Two forward-power couplers were built into the RF output line of the tube. One was positioned ahead of the output window, on the vacuum side, and the other was placed on the output side of the window. If the signals from both couplers were




Figure 3-5. Schematic diagram of acoustic waveguide-arc detector.

not nearly identical, a window arc was inferred.

3.16 Forward and reverse couplers

It is at this point in the power-transport system that the performance of the individual power amplifiers can best be assessed. Here, reverse-power couplers give a full-time indication of the state of the remainder of the high-power transmission system, which may include rotary joints, radomes, switches, and other devices that can wear out, get water-logged, or get stuck halfway if they are two-position devices. The couplers also give early warning for arcs that may have started downstream and that, if permitted to persist, will propagate toward the source of power at the speed of sound.

3.17 High-power circulator

Whether or not to put a high-power circulator between a microwave power tube and its output load is one of the most agonizing decision that any transmitter designer must make. Tube manufacturers will make it easy for you. They usually insist that you use one or their warranty will be void. (The same thing goes for the arc detector and the electronic crowbar.) The decision, nevertheless, is the designer's.

A high-power circulator has only three traits that a designer can absolutely depend on:

- it will be one of the most expensive components in the system, usually second only to the tube itself;
- it will have one of the longest lead times; and
- it will absorb precious transmitter output power and convert it to heat, which must be carefully removed to avoid damaging the circulator.

What will it do in return? It will absorb reflected power from the load system. That's all.

There is no reason why a properly designed load system should have a voltage/standing-wave ratio (VSWR) greater than 1.5:1, and there is no reason why a properly designed microwave power tube should not tolerate a VSWR of 1.5:1. If, for whatever reason, either parameter cannot be realized, then and only then should the subject of the high-power circulator be broached—at which time it may be discovered that there are no credible high-power circulator sources, anyway. (It has been said that what makes a manufacturer a source of high-power circulators is the glue that he uses to fasten the ferrite material to the waveguide wall.)

3.18 Output hybrid combiner

In the system that we are describing, the output hybrid combiner is a short-slot, 90°, 3-dB hybrid coupler. The twin of (3.8), the splitter, it is identical to it in all respects, except for the tailoring of its edges to minimize electric-field enhancement. This is done in the interest of maximizing power-handling capability.

It should be mentioned again that it is not necessary for the outputs of the HPAs to be equal for them to be combined with 100% combining efficiency, or

that the number of HPAs must follow a 2^n progression. (Later, a transmitter comprising 24 high-power TWTs whose outputs were combined in a single-output waveguide will be described.)

3.19 Waster power coupler

The output of this signal sampler most effectively indicates how well the outputs of the two power amplifiers are being combined. This measurement can best be taken when the phase shifter (3.10) and the time-delay shim (3.11) have been adjusted to minimize the frequency-averaged amplitude of this signal over the operating bandwidth of the system. Presuming that the reflection coefficient of the waster load is small (as it should be), the directivity of this coupler is not of primary importance.

3.20 Waster load

Assuming there is proper input phasing and reasonable amplitude balance, the nominal power dissipated in this load will be but a small fraction of the total combined RF power. Nevertheless, gross maladjustment of the phasing can put the total combined RF power into this load instead of the nominal output port of the combiner. In more than a few systems this effect has been exploited to obviate the necessity of placing a high-power switch between useful and dummy output loads, in which case this load is designed to perform accordingly. Where the waster load is sized to handle only the anticipated waster power (with a prudent safety factor), the coupler (3.19) can be set up to shut off the RF drive to the transmitter if power in this line is above a safe threshold. The sudden loss of power output from one or the other HPAs will cause half of the power of the remaining one—or 1/4 of the total combined RF power—to be diverted to the waster load.

3.21 Harmonic filter

Among transmitter designers, placing a harmonic filter in the waveguide system is almost as controversial as including a high-power ferrite circulator. In situations where it is known that harmonics of the carrier frequency will coincide with known operating frequencies of other nearby RF facilities, the designer has no choice: the harmonics must be attenuated to the greatest extent possible. Moreover, the transmitter must absorb the harmonic components. Reflecting them may keep them from being radiated, but in the worst case they can still produce "trapped" undamped harmonic resonances within the waveguide system. The presence of such resonances in the waveguide system will cause RF breakdown even at a fraction of full power. Furthermore, in dominant-mode waveguide systems all harmonics are high enough in frequency to possibly cause the propagation of higher-order modes. And the higher the harmonic number, the larger the number of possible modes.

This can be either good or bad. If the problem is presented in terms of a maximum permissible level of total harmonic power in the waveguide system, then the filter must be capable of absorbing the power being propagated in all of the possible modes. If, on the other hand, the designer is only concerned with what is actually being radiated into space, the antenna and its feed structure may act as an effective mode filter, keeping harmonic power from being radiated at

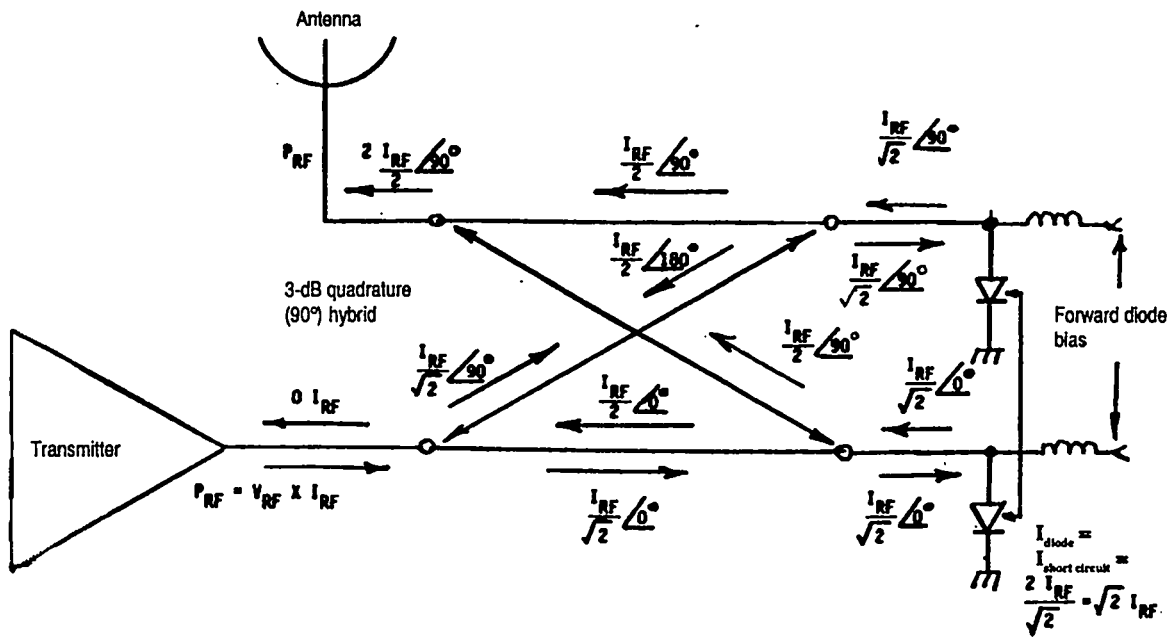


Figure 3-6. Vector relationships in the diode duplexer for RF current in transmit mode.

all. So-called leaky-wall filters have multi-mode absorptive capability. So do the “waffle” type filters. All filters will have some impact on the fundamental-frequency power-handling capability of the waveguide. Most of the time they reduce it. But if they are deployed ahead of line sections with harmonic resonances, they can actually improve the performance of the total system. Where there are specific interference problems that involve a high-order harmonic (such as the fifth or seventh), special traps can be used. In such traps, the miter wall of a right-angle bend is not solid but is reflective enough at the fundamental frequency so that the fundamental can negotiate the bend, whereas high-order harmonics pass through the virtual miter wall into an absorber.

3.22 Antenna duplexer

The microwave component depicted in Fig. 3-1 as the antenna duplexer is used in monostatic radar systems to permit time-sharing of a single antenna by both transmitter and receiver. There are two basic forms of such devices. One is called a branch duplexer and the other, the balanced duplexer.

The one shown in Fig. 3-1 is the balanced type, comprising high- and low-power, 3dB, 90° hybrid junctions and the equivalent of RF shorting switches. The transmission lines from the transmitter to the antenna are connected to adjacent ports of the high-power hybrid, whose output ports are also shunted by the high-speed RF shorting switches. Gas-filled cells are commonly used as the shorting switches. In them gas is ionized by incident transmitter power and the resulting plasma short-circuits the junctions, reflecting the incident transmitter power back through the hybrid.

Because of the 90°-phase shifts in both directions, the reflected signal components combine in the antenna arm and cancel at the transmitter arm, effectively connecting the transmitter to the antenna. Following each transmitter pulse, the gas de-ionizes, eventually returning to an open-circuit state (over a period called




Figure 3-7. Vector relationships in dipole duplexer for RF current in receive mode.

the recovery time). At this time, the returning radar signals reflected from the target (called returns) travel from the antenna to the antenna port of the high-power hybrid. Here, they split. But this time they continue past the high-power hybrid output ports, which are no longer short-circuited by the ionized gas, to the low-power hybrid input ports. The receiver input is connected to one output port and a waster termination to the other. Assuming perfect amplitude split and exact 90° -phase shift in the two hybrids, all of the target return power will combine at the receiver input, which is diagonally opposite the high-power antenna terminal. Conversely, during the transmit pulse, signals that leak past the intended short-circuit impedance cancel at the receiver input port and combine at the waster-load port, which diagonally opposes the transmitter input port.

Unfortunately, gas cells absorb transmitter power in the process of ionization, which is not instantaneous. During this time there is a short-duration leakage, called spike leakage, past the gas cells. If all goes as intended, all of this leakage ends up in the low-power hybrid waster load. Often, however, enough of it does not end up there. In such cases, the unwanted power can burn out sensitive low-noise amplifiers in the receiver.

Solid-state PIN diodes can be made to exhibit low values of RF impedance when biased with current in the forward direction and relatively high values of RF impedance when biased with voltage in the reverse direction. Especially at the lower microwave frequencies, arrays of such diodes have been used as the RF switches in very high-power duplexers. As switches, they have the disadvantage of having to be actively biased. Even so, this presents the opportunity to start the forward current bias before the leading edge of the transmitter pulse, eliminating spike leakage altogether. Only the so-called flat leakage, which persists throughout the transmitter pulse, remains.

Figure 3-6 shows the vector relationships for RF transmitter output current when the duplexer diodes are forward-biased in the transmit mode. And Fig. 3-7 shows the same relationship for received-signal RF current when the duplexer diodes are reverse-biased in the receive mode. The relationships shown in Fig. 3-6 also hold true for reflections from the RF power amplifier tubes with regard to the input hybrid power splitter (3.8) and for reflections from the RF power ampli-

fier output terminals with regard to the output hybrid power combiner (3.18). The vector relationships shown in Fig. 3-7 hold true for the splitting and recombining process involving the high-power RF amplifiers, except that the diode junctions are replaced with identical amplifiers so that the vectors to the right of the junctions are greater by the current gain of the amplifiers than those to the left.

Figure 3-8 shows the dimensions and arrangement of a practical multi-megawatt solid-state duplexer using PIN diode switches. The high-power hybrid junction is in a WR-2100 rectangular waveguide. Its two output ports are integrated with cross-bar-type waveguide/coaxial transitions to 9-1/8-in. rigid coaxial transmission line. Connected to each transition is a PIN diode stub, the details of which are shown in Fig. 3-9. The stubs are built into quick-step transitions between 9-1/8-in. and 7/8-in. rigid coaxial line. The 10 diodes in each stub are arranged in a circle concentric with the center line of the coaxial transmission lines, bridging that portion of the transition region where the electric and magnetic fields are radial rather than coaxial. The PIN diodes themselves, are so small that their total length occupies only a small fraction of the total height of the radial section. The remainder of the height is made up by the screw-in diode holders, which allow individual diodes to be removed without disassembly (other than removing the weather-proof top cover). Even at the relatively low operating frequency of 425 MHz, the capacitive susceptance of the back-biased diodes is so




Figure 3-8. 30-megawatt UHF solid-state diode duplexer.

great that a very large mismatch occurs at their insertion radius. For that reason, the bias-insertion leads, which alternate circumferentially with the diode holders, have a length and a shape appropriate for the inductance required for parallel resonance with each diode. The forward and reverse diode bias signals are isolated from the UHF signal path by a segmented coaxial blocking capacitor. (This coaxial transmission line section has very low characteristic impedance and uses thin-wall Teflon tubing as dielectric.) The outer conductor is broken into 10 circumferential segments so that individual diode bias conditions can be monitored.

Connected to the diode stubs by 7/8-in. coaxial line is the low-power hybrid, the output ports of which connect to the receiver input and the waster load.

The design shown in Fig. 3-8 was tested for 30 MW peak power, 15- μ s pulse duration, and 0.005 duty factor. It operated for many years at 20 MW peak power. Because of the small mass and thermal time-constant of the PIN junctions, the stress that limits diode performance is intrapulse temperature rise, which is proportional to individual pulse energy (the product of peak power and pulse duration).

Before leaving the subject of duplexing, however, it should be contrasted with the process of diplexing. According to the dictionary, both of these terms mean substantially the same thing. In RF system usage, the term "to duplex" means to separate functions in time; "to diplex" means to separate functions by frequency or wavelength. A common example of high-level diplexing is the sharing of a common antenna by the video and audio transmitters of a typical high-power television transmitter. The main signal is a single-sideband AM picture signal and the other a narrow-band (25-kHz deviation) FM sound signal, with carrier frequencies separated by 3.5 MHz. Both signals are coupled into the antenna through filters—one low-pass and the other high-pass—that function much like the cross-over network used in two-way loudspeaker systems.

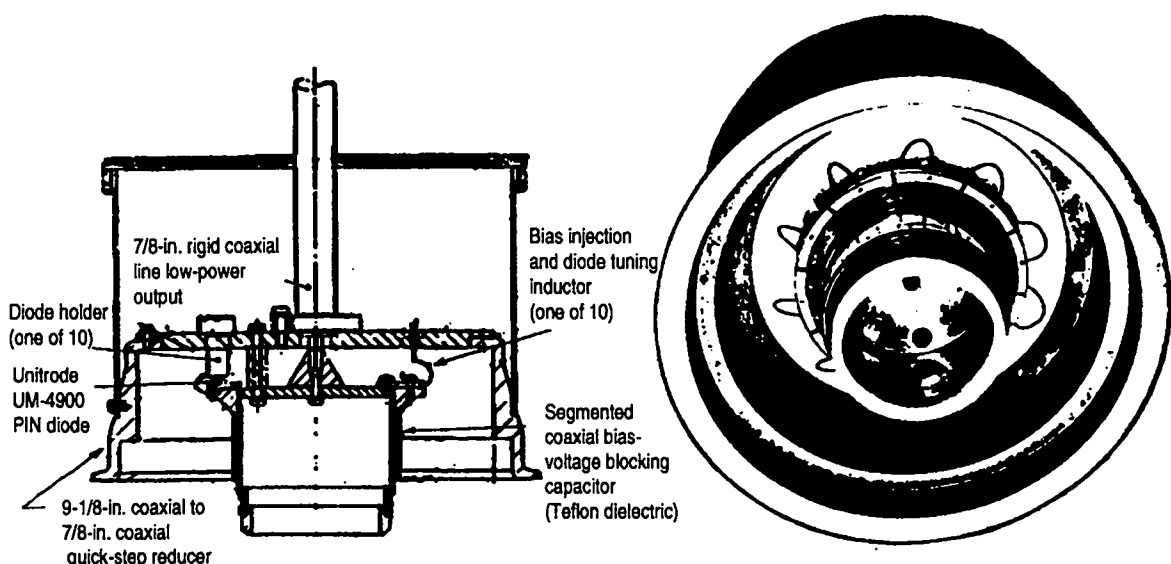


Figure 3-9. Details of 30-MW diode duplexer diode stub.

3.23 Waveguide switch

In many high-power RF systems it is necessary or desirable to be able to switch the transmitter output from the antenna to a dummy load. As mentioned before, this effect can be achieved by deliberately mis-phasing the two amplification channels so that most of the transmitter power combines in the hybrid combiner waster load. The key word is "most." There will always be a significant residue directed to the antenna—more, usually, than can be tolerated. The isolation of a mechanical waveguide or coaxial switch can be made much greater by using a waveguide switch. Moreover, the antenna port can be covered by a metal shutter as well.

3.24 High-power dummy load

At this location, a full-power dummy load (or "phantom antenna," as they were once called) can permit full-power transmitter operational diagnostics even when the antenna is unavailable or radiation is not permitted. If the load is water cooled, or is in fact a true water load where the water is the lossy medium as well as the coolant, it can be set up as a calorimeter in order to yield what some maintain is the most accurate measurement of average power output. (Yet this is another controversial subject).

4. Transmitter Engineering and the 'Iilities'

Fear and loathing on the transmitter-design trail.

No discussion of high-power transmitter design would be complete without mention of the impact of the "ilities," of which the most familiar are

- Reliability,
- Maintainability,
- Availability,
- Electromagnetic susceptibility, and
- Electromagnetic compatibility.

In the same category, but not literally "ilities," are

- Safety,
- Built-in test equipment,
- Failure-mode and effects analysis,
- Human factors, and
- Configuration control.

In the full realization that entire books have been written on each of the above subjects—and that the subjects themselves produce fear and loathing in most red-blooded design engineers—the author shall nevertheless hazard a few comments.

4.1 Reliability

Most engineers, especially those who have worked on military contracts, are familiar with, or at least have heard of, MIL-HDBK 217, the so-called reliability bible. Pentagon officials would like to think that an engineer with MIL-HDBK 217 in one hand and an electronics parts list in the other should be able to calculate any electronic system's reliability in terms of the mean time between failure (MTBF). Except for certain parallel-redundant, fault-tolerant special circuits, the military assumes that the failure of a single component should result in the failure of the circuit to perform properly. (After all, if the circuit performs properly without the component, it wasn't needed in the first place, right?) In which case, the total failure rate of the circuit is the sum of the failure rates of the individual components. These individual failure rates can be found in MIL-HDBK 217. The MTBF is the reciprocal of the total failure rate, except for one thing. The failure rates in the handbook are *random* failure rates. These are the rates that apply throughout the flat bottom of the so-called bathtub-shaped curve that describes the useful life of a component. The downward slope at the beginning of the bathtub is due to "infant mortality" of the less robust components, and the tilt up at the other end is indicative of the "end-of-life syndrome." Unfortunately, these statistical preconceptions do not always apply to transmitters.

In the first place, for high-power transmitters few components have the luxury of random failure. Systematic and catastrophic failures are more the rule than

the exception. Stress levels have many more dimensions than those that can be encompassed by a reliability handbook. For instance, take the case of a conventional wire-wound resistor performing the function of current-limiting resistance described earlier in Chapter 2, Section 6. This resistor will have average-power dissipation because the flow of pulse current through it is on a repetitive basis. Its thermal properties and/or external cooling must be adequate to limit its average temperature. But there are other considerations. There will be pulse heating in the resistor because the current flowing through it is in the form of pulses. Worse yet, if there is a high-voltage arc in the system downstream of the resistor, perhaps all of the energy stored in the system will be dissipated in the resistor in so short a time that heat will not be able to leave it. In addition, all of the system voltage will be applied across the resistor. And that is not the end of it. The fault current (which, indeed, the resistor is intended to limit) may be 50 to 100 times as great as the normal pulse current. Finally, magnetic pressure, which is proportional to the square of this current, will attempt to unravel the resistor, and this force is by no means negligible. In fact, all components in the high-voltage/high-current path are subjected to similar sets of layered stresses.

Even low-level components, which seemingly should be farther out of harm's way, can be subjected to induced, convected, and conducted impulse stresses that can produce instantaneous mortality or, sometimes even worse, weakened or altered states. In modern transmitter designs (except those designed by the Russians), low-level active components are exclusively solid-state. Easily 10% of the total component count in a well-thought-out design can be devoted to circuit elements whose purpose is to protect prime-mission components from damage or upset. Especially susceptible to damage are an integrated circuit's gate inputs, which must be current-limited and diode-clamped for survival. If you think that using solid-state devices with large-area junctions will eliminate the problem, think again. Solid-state junctions having a large surface area do not, in themselves, assure survivability, especially when subjected to short-duration overvoltages. Solid-state junctions fail because they are short-circuited by the growth of metallic lattices from one side of the junction to the other, a process as inexorable as the loss of emissive material from a thermionic cathode. As mentioned earlier, the rate of this chemical reaction is governed by the Arrhenius equation, which states that the rate is a log-normal function of temperature. (This is why the failure-rate data given in MIL-HDBK 217 is for solid-state junctions whose temperature remains *below* 175°C.)

Solid-state junctions do not last forever. When they are subjected to short-duration overvoltage, current does not have time to spread out over the full surface area of the junction, nor does heat have time to be conducted away. An extremely high temperature, therefore, may be confined to a very small region of the total junction. The effect of this condition would be analogous to a large number of small junctions wired in parallel; the failure of one is all that is required for device failure.

There are subordinate attributes of devices that can be vehicles of failure as well. The power MOSFET, a relatively high-voltage, high-current switch that is used in many of the low-level transmitter circuits, has a residual bipolar transistor that manifests itself as an anti-parallel diode. Voltage reversal caused by

abnormal external circuit conditions will produce current flow in this diode. If the flow is too much or too fast, the diode will fail, and the source-drain connections will be short-circuited. All solid-state diodes, after conducting in the forward direction, will conduct in the reverse direction even more vigorously until the charge has been removed. If this current is not limited, even "fast-recovery" diodes can be destroyed. (Later, we will describe two solid-state diode applications that had to be retrofitted in every sense of the word with thermionic, or "real," diodes because of the reverse-conduction effect—and ironically they had to be retro-designed *back* to solid-state devices because it was discovered that nobody manufactures high-voltage thermionic diodes any more!)

Multilayered stress affects almost all electronic components. That is why it is not really practical to rely on such guidebooks as MIL-HDBK 217, which only recount single-point failure rates and tell nothing about systematic or catastrophic failure rates. Designers must take into account the complicated array of potential stresses present in the system.

4.2 Maintainability

If the prospects for designing the perfectly reliable transmitter are bleak, it would be a good idea to devote some engineering attention to the maintainability of a transmitter. Fortunately, this is easier to do: design the transmitter so that all its components can be easily and quickly accessed for repair and troubleshooting.

One tenet of Murphy's Law states that the component most likely to fail is the one that is least accessible. (An as-yet-unproven corollary states that if all components are equally accessible, none will fail.) There are only two guidelines for designing transmitters with good maintainability:

1. Never mount a component on another component; and
2. Never mount a component or subassembly in such a place that you have to remove another one to get to it.

Needless to say, these guidelines are violated all of the time, and for perfectly justifiable reasons. But designers should at least be aware that they are violating a guideline. (A frequently encountered barrier to maintainability is the hardware used to affix covers and access plates. Every screw should be justified. If they are all needed for structural reasons or to keep wanted things in and unwanted things out, the fasteners had better be captive. If they aren't, you can bet that they will never be completely replaced once they have been removed.)

4.3 Availability

Availability is the bottom-line combination of reliability and maintainability. In most cases, it is what really matters to a user or customer. For instance, if a transmitter fails once per year on average and takes eight hours to repair on average, it is unavailable for service eight out of every 10,000 hours (roughly). Which means that it is available 9992 hours out of 10,000 hours, or 99.92% of the time. (This same availability could be achieved if the transmitter failed more

frequently but could be repaired more rapidly.) Designing for maximum reliability is laudable and sometimes mandated, but is not always cost-effective or even affordable. Designing for ease and speed of maintenance is always cost-effective, especially if your reliability-enhancement efforts have done nothing more than identify those components that are most likely to fail.

Among the most costly transmitter designs are those with "100% availability." These are the parallel-redundant, multi-channel transmitters that possess something described as "graceful degradation." They have 100% availability, but at something less than 100% of maximum capability. All high-power transmitters with solid-state RF power amplifiers fall into this category, simply because it is not possible to achieve high power from solid-state sources except by the parallel combination of large numbers of individual modules.

But multiple-channel transmitters using microwave vacuum-tube power amplifiers have also been designed. The two biggest, the PARCS (Perimeter Acquisition Radar Characterization System), and the Cobra Dane radar, use multiple high-power TWTs and are phased-array systems. The PARCS uses 128 Raytheon PPA-200 UHF TWTs configured as 16 redundant eight-tube transmitters. The Cobra Dane transmitter uses 96 Raytheon L-band TWTs configured as 12 redundant eight-tube transmitters. The outputs of the TWTs are combined in space, as each tube feeds multiple antenna elements through phase-shifting devices. The loss of a single tube reduces the radiated power to $(N-1)/N$, but reduces the far-field power density to $(N-1)^2/N^2$. (This calculation takes into account not only the lost power but the lost antenna aperture as well.) In a transmitter such as the ALTAIR UHF transmitter, which combines the outputs of three PARCS-type eight-TWT transmitter subsystems into a single-waveguide antenna feed line, the loss of a single TWT reduces output to $(23/24)^2$ of the maximum total. An amount equal to $1/24$ of the total is, obviously, the direct loss due to the defunct tube. A less obvious amount equal to almost another $1/24$ results from an amplitude imbalance component of waster-load power throughout the power combiner matrix. Transmitters of this type degrade gracefully, but not quite as gracefully as their most ardent proponents would have us believe.

4.4 Electromagnetic susceptibility

How susceptible are transmitter circuits to upset or even destruction resulting from self-generated or external electromagnetic events? A designer may have no way of knowing until it's too late. The most tragic example of the effects of electromagnetic susceptibility was the fly-by-wire military helicopter, which lost its computer-generated control when it flew too close to certain radar systems and crashed (both figuratively and literally). All electronic circuits—digital, analog, or RF—are susceptible to external RF interference if it is at a sufficient power level and at an offensive frequency. For instance, there are few computers or word processors on the small Pacific island of Roi-Namur that can ignore the ALTAIR 7-MW, 162-MHz VHF transmitter when its 150-ft-diameter dish is aimed at the horizon and pointed in their direction. This is because transistor-transistor logic (TTL) circuits in the equipment depend upon the very small difference between the collector-saturation voltage of one transistor and the base-emitter turn-on voltage of the succeeding one to differentiate between a logical 0 and

logical 1. The noise immunity designed into these circuits is virtually nil. Modern computer circuits are also fast enough to respond directly to bipolar-induced VHF signals. Hence, there is a guaranteed instant upset.

Yet the timing synchronization circuits for the ALTAIR, which share the same building as the transmitter, are oblivious to the same radiation. Still, they use integrated circuits that are functionally similar to the ones on Roi-Namur. What makes them different is that successive transistors are Zener-diode coupled. This feature introduces noise immunity equal to the Zener diode voltage. It also makes their gate response slower, but it is still more than fast enough for their dedicated purpose. The family of logic elements used for ALTAIR is Motorola's MHTL (Motorola high-threshold logic), which is only available for replacement purposes and is not recommended for new circuit designs. Teledyne makes a similar family called HINIL (high noise-immunity logic), and it serves the same purpose. (Members of this family have found their way into the antenna-servo logic of the same radar and for the same reason.)

In short, do not underestimate the impact of electromagnetic susceptibility on your design. As a minimum, designs *must* tolerate the radiation from the RF power source they are intended to serve. (It's also a good idea, as it certainly was on the island of Roi-Namur, to know who your neighbors are.)

4.5 Electromagnetic compatibility

The mirror-image of electromagnetic susceptibility is electromagnetic compatibility. To what extent will one piece of equipment intrude on another? When requirements for electromagnetic compatibility are expressly specified, they usually apply to so-called spurious or harmonic emanations, those that are superfluous to the prime mission of the equipment but that can be quite troublesome to adjacent equipment. For example, on the island of Roi-Namur there are high-power radar systems at VHF, UHF, L-band, S-band, C-band, K_A-band, and W-band—all but the last two of which are in the megawatt or multi-megawatt class. Not only must all of the transmitters—except the last two—meet stringent harmonic-suppression specifications, but the fundamental operating frequencies and band edges must be coordinated so that the harmonics of one do not coincide with or intrude into the operating band of another.

But the real issue here is interference. Put as simply as possible, if a system has excessive electromagnetic susceptibility, the designer had better head back to the drawing board, because even a superb job of engineering compatibility on the part of a transmitter's electromagnetic neighbors will not make that transmitter immune to their fundamental-frequency emissions. (It is worthy of note, however, that in some extremely crowded electromagnetic environments, fundamental-frequency modulation envelopes have been mandated. One of the most difficult to synthesize, but least sideband-intrusive, is the Gaussian envelope, which has the unique property of having identical shapes in both the time and frequency domains).

4.6 Safety

There are those who insist that designing a transmitter for safety is almost as important as designing it for its primary performance goals. They are wrong. It




Figure 4-1. An example of a mechanically "cheatable" interlock switch.

is *more* important.

It goes without saying that all microwave vacuum-tube transmitters require operating voltages that can cause severe injury or death to a person who comes into contact with them. How high do these voltages have to be? Military specifications mandate that anything over 30 V be considered hazardous (which means that even all-solid-state transmitters are not exempt from safety-design considerations).

Exposure to hazardous voltages must be prevented by enclosing their sources. Access panels or doors to these areas must be interlocked in such a way that opening them automatically removes the hazardous voltage *and* short-circuits any circuit element capable of storing a charge or energy, even after the source has been interrupted. Electrical interlock switches are made especially for this purpose. Some have mechanical means whereby they may be "cheated" so that internal electrical power can be restored even with the access door or panel open. An example of a commonly used door switch is shown in Fig. 4-1. Note that the switch is normally actuated when the plunger is pushed in, as when a door or panel is closed or a lid is let down. When the door is opened, the spring deactuates the switch, thus opening the circuit. However, if the plunger is pulled farther forward by an impatient technician, the detent is overridden in the opposite direction and the switch is once again actuated. Unfortunately, this type of interlock is often used unwisely (or perhaps even illegally in this day of product liability). Where voltages are above 300 V, military standards forbid the use of "cheatable" interlock switches. (It goes almost without saying that any form of interlock system can be "cheated" if a technician has an adequate supply of time, tools, jumper wire, and nerve.) This complaint refers only to the type of switch shown in Fig. 4-1, which can be easily defeated by an unassisted individual. An especially robust form of "uncheatable" interlock is a mating male/female connector pair, as illustrated in Fig. 4-2.

What is important in designing for safety? It is *obvious* simplicity; it is defi-

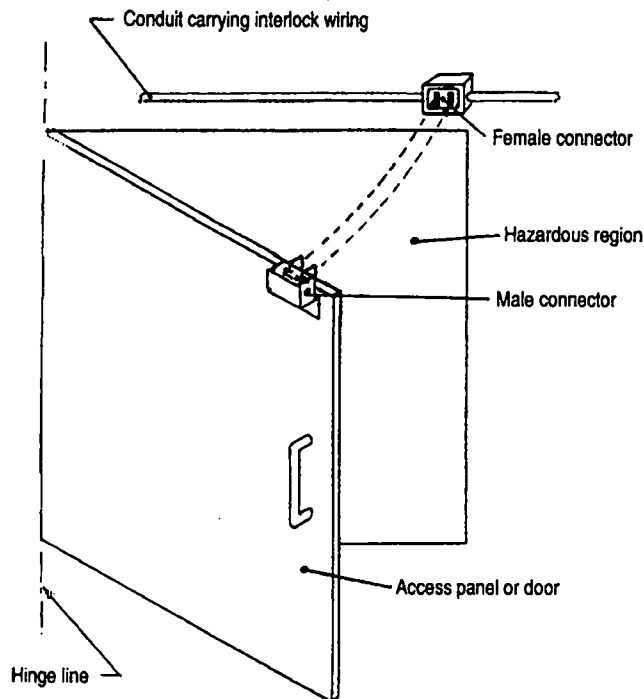


Figure 4-2. An absurdly simple, but surprisingly effective, non-cheatable interlock.

nitely *not* sophistication. When workers breach an interlocked enclosure, it is critical for them to know right away that the electrical de-energizing function of the interlock is working. An example of such a no-nonsense interlock can be seen in Fig. 4-2. In this switch, the continuity of the electrical circuit requires that the male connector, whose pins are jumpered together, be inserted in the female connector, one of whose terminals is connected to the power source while the other is connected to the load. The male connector is attached to the door, the female to the door frame. When the door is opened, the male jumper is pulled out of the female connector, thus opening the circuit. The status of the switch, and therefore the circuit, is immediately obvious to the person opening the door. This is a good interlock.

On the other hand, had the switch shown in Fig. 4-1 been used, the technician would never really know for sure if the interlock was working. In this case, the opening of the door removed the pressure that had kept an electrical switch in the closed position. The fact that the mechanical actuator of the switch had changed position may be obvious to the worker, but it would be by no means obvious that the electrical circuit has been broken. The key word is obvious. For interlock switches, the more obvious, the better.

The opening of an electrical-interlock circuit must also positively remove the hazard. Examples of such action are open-circuiting the hazardous-voltage conductors at the source end or short-circuiting the hazardous-voltage conductors at the load end. Preferably both should occur, in the order stated. One typical high-voltage shorting electro-mechanical relay is shown in Fig. 4-3. It is normally in the short-circuit position as shown; its solenoid electromagnet must be energized in order for it to hold its shorting bar in the raised position. If for any

reason power is interrupted to the switch, the circuit is automatically shorted. It is, therefore, at least electrically fail-safe. A good designer would typically install such a relay in parallel with the transmitter's electronic crowbar (see Fig. 2-1, item 6).

Open-circuiting the hazardous voltage conductors at the source end can be accomplished with a series-opening switch. Figure 2-1 indicates two possible locations for it: the main circuit breaker of the unit substation (13) and the high-speed contactor (12). A favorite ploy of many safety-oriented designers is to make this switch the main circuit breaker, as long as it has an undervoltage trip feature. This feature, which is a field-installable option in some of the larger frame sizes, is often confused with the shunt-trip feature. It is the latter's logical complement. Whereas the shunt trip requires the presence of an external voltage to cause the breaker to trip, the undervoltage trip requires the presence of an externally supplied voltage in order to stay closed. Should this voltage be removed (or fall below a predetermined threshold), the breaker will automatically trip, and it cannot be reclosed until the external voltage has been reapplied. This feature makes it compatible with a fail- and connectivity-safe interlock topology and is highly recommended.

Before moving on to key-based interlocks, one more thing needs to be said about electronic crowbars. Although a crowbar can never be a proper component of a safety-interlock system, it has demonstrated more than once its capability to be a genuine life-saver—but only as the microsecond-response component




Figure 4-3. A typical high-voltage, electrically operated shorting relay.

of a ground-fault-interrupt (GFI) system involving the design of the entire high-voltage circuit. Harking back to Chapter 2, Section 10, where the high-voltage circuit was likened to household electrical service—with black (high-voltage), white (load-current return) and green (safety ground) conductors—Fig. 4-4 shows how this can apply to personnel lifesaving in exactly the same fashion as a bathroom GFI circuit. In the figure, the load on the high-voltage system is a microwave vacuum tube with an insulated (or isolated) beam collector. The rest of the tube envelope is metal and is fastened to ground, which establishes the single-point ground for the system. Both terminals of the high-voltage power supply and energy-storage capacitor are insulated from ground (black and white wires).

All current that flows in ground must also flow through the current-monitor transformer, which surrounds the single conductor that connects the low-voltage return conductor from the power supply to ground at or near the tube envelope. Assuming that the electron beam of the tube is well focused, about 99% of the beam current will flow directly between cathode and collector (dashed line) and only 1% will be intercepted by the ground part of the tube, called the body. Therefore, the normal current through the current monitor, even for a very high-power tube, may be less than one ampere. The threshold for the low-level firing circuit for the electronic crowbar, then, can be set for something only slightly greater.

The stage has now been set for our hero, who, standing on ground, has come in contact with the power supply high-voltage lead. Current through his (or her)




Figure 4-4. The electronic crowbar as a safety ground-fault interrupter.

body is treated no differently than microwave tube body current, and as soon as it exceeds the crowbar threshold, the crowbar will fire, diverting the lethal charge from our hero within microseconds. Actual lives have been saved in this way. (It is worth remembering that the French term for tube body current is I_{corpse} .) In order to be fully effective, however, the majority of the total current in the high-voltage loop must flow in a current path that does not include ground.

Automatic electrical interlocking is sometimes augmented by a key-operated system, the most common one being the "Kirk" key system, which can become quite elaborate. A typical system might involve a master key-operated switch, which shuts down power to a console or locks out a manually operated power-disconnect switch by extending a bolt. (The mating pieces of the system are pin-keyed so that the bolt cannot be extended unless the subassembly that includes the bolt is engaged with the piece that receives the bolt to effect the lockout.) Once the lockout has been accomplished, the key can be removed from the lock. (Until then, the key is trapped in the lock and cannot be removed.) The key, once removed, can now be used in a complementary subassembly to withdraw a bolt that may, for instance, have been preventing the opening of the door to a hazardous region. So long as this door is open, the key is again trapped, and can only be removed when the door has been closed and the bolt extended by rotating the key in the lock.

If there are multiple areas that may require simultaneous access, a component called a key-transfer assembly can be used. In theory, any number of keys can be trapped in such an assembly. A key can be freed only when the single access key is inserted in the proper lock and rotated. This procedure, in turn, traps the access key in the lock. Any or all of the subordinate keys may then be removed. These subordinate keys can be used to unlock as many areas as there are keys. The master access key, however, cannot be removed from the transfer block until *all* of the subordinate keys have been returned to their places in the block.

In order to assure safety, all of the keys must be coordinated and no duplicates can be tolerated (even in another system). It should be obvious by now, however, that the loss or misplacement of a single key would mean that the system could not be turned on. For this reason, the operators of most high-priority systems—and what high-power RF system would not consider itself high-priority?—have managed, one way or another, to obtain duplicate keys. (They insist the keys are kept in a "secure" place and under tight administrative control. Nevertheless, the entire safeguard philosophy of the system is subverted by this act.)

The hazards of high-power, high-voltage systems do not end with the possibility of electrocution, however. Stored electrical energy can be released with explosive mechanical force in the form of shattered solid-dielectric insulating bushings, the shards of which can have the same lethality as shrapnel. It has been calculated that one megajoule of stored electrical energy has the same explosive potential as half a pound of TNT. There is no reason to doubt this. The best enclosure for a large capacitor bank should act as both a shielded room and a blast vault.

Workers on high-voltage systems have also been known to drown or be asphyxiated as well. Systems operating at voltages much above 50 kV are often

insulated with dielectric oil that is either petroleum- or silicone-based. These oils are less dense than water and often fill containment vessels that rival small swimming pools in volume. Therefore, they take considerable time to drain and refill. It is tempting to minimize repair or maintenance time by entering the tank without draining the oil. *Some* people have negative buoyancy in water. *Everyone* has negative buoyancy in oil. And a few have drowned because of it. Sulfur hexafluoride (SF_6) is also frequently used as a dielectric fluid. It is a colorless, odorless gas that is heavier than air, and it can fill a tank in exactly the same fashion as oil, displacing all of the air that originally filled the tank. If the tank is large enough for someone to enter it and become totally immersed in the SF_6 , he or she can become asphyxiated without any warning from the body's defense mechanisms. At least one such death has been attributed to this cause. Large volumes of SF_6 are difficult, if not impossible, to make safe. They should simply be avoided in the design stage, if at all possible.

There are also two physiological hazards from RF radiation: the "microwave-oven" (or "English-muffin") effect and the acoustic-shock effect. The microwave-oven effect is the best known and understood. It involves the heating of internal organs and tissues when flesh and blood are exposed to microwave radiation. The acoustic-shock effect is caused by the kinetic energy contained in sonic waves that can emanate from the transmitter. These waves are especially hazardous to eyes, but the effect is only present in pulse-modulated RF systems.

High-power microwave-tube systems also operate at voltages high enough to generate ionizing radiation in the form of x-rays, which must be dealt with (usually in the form of added shielding). The latest, and most controversial, of the possible electrical hazards is extremely low-frequency (ELF) magnetic fields (in the milligauss range). Included in the range of suspect frequencies are the most popular power-line frequencies, 50 and 60 Hz, and some popular radar-system pulse-repetition rates, the most common of which is 30 pulses per second (pps). Exposure to high-current, high-voltage ELF has been linked in some studies to the development of leukemia in children. It would behoove the circumspect transmitter designer to stay abreast of the scientific and popular literature concerning ELF magnetic fields.

For better or worse, safety and safe design practices are increasingly becoming institutionalized. Environment, safety, and health organizations are becoming more pervasive and influential in government and private industry, gaining the administrative power to shut down systems that violate safety regulations. In the future, no high-power system design will be considered adequate unless it treats safety with the same, or greater, concern as its prime mission objectives.

4.7 Built-in test equipment

Just to insure its own proper functioning, any adequately designed high-power transmitter monitor and control system will inherently contain more test-point sampling points than most external users of such data are prepared to handle. If this is not the case, it is probably time to review the efficacy of the control and monitor system. (Later, as specific designs are discussed, the role of built-in test equipment will become clearer.)

4.8 Failure modes and effects analysis

In any high-power engineering project, the analysis of failure modes and their effects is a serious matter. The failure of a key component to function properly can initiate a true "domino" effect, leading to extensive subordinate (or collateral) component failure—or even to the electronic equivalent of meltdown. A memorable example of this took place at Dallas-based Continental Electronics during the testing of a 1-MW broadcast-band transmitter built for the Saudi Arabian government. The giant tetrode output tubes were protected from internal arc damage by an electronic crowbar in the form of an ignitron. When it fired for the first time, short-circuiting the entire high-voltage system preceding it, the special-purpose high-speed primary switch gear that was designed to disconnect the faulted high-voltage system from the primary ac mains not only did *not* respond with the intended speed, it did not respond at all! The faulted high-voltage system thus remained connected to East Dallas Power and Light. All the lights in the building dimmed until the valiant ignitron had been reduced to molten copper and a puddle of mercury. The high-voltage transformer/rectifier survived, thanks to a conservative design. What could have been a disaster resulted mostly in a spectacular short-term power bill. (Later it was discovered that a well-intentioned but not-so-well-informed technician found some "loose" hardware in the switch gear and tightened it down, thus preventing it from actuating properly.)

This little anecdote represents a good-news, bad-news story for Continental Electronics' failure-mode analysis. The good news was that the system's major components were designed to be robust enough to endure such an insult and suffer little or no damage. The bad news was that somebody forgot to predict the domino effect if the primary switch gear failed.

4.9 Human factors

Since 1949 the term for human engineering has been "ergonomics," which, as Webster's Ninth New Collegiate Dictionary states, is "an applied science concerned with the characteristics of people that need to be considered in designing and arranging things that they use in order that people and things will interact most effectively and safely." In this field, there are specialists who often have great awareness of the capabilities and limitations of humans in their interaction with machines. However, a properly designed transmitter, as will be discussed later under the subject of control and monitor, will act autonomously with electronic speed in order to safeguard itself from hazardous conditions. In some situations, there simply isn't time for human intervention. On the other hand, the control screen or control-panel displays, which alert the human operator as to what has automatically transpired and the reasons for it, must tell the story unambiguously.

Except for the most recent of transmitter control-system designs, such communication usually takes the form of indicator lamps that in the simplest case were either illuminated to show a function was enabled or off to show that it wasn't. The obvious source of ambiguity in this case is the burned-out lamp. More sophisticated designs used multicolored lamps: green for OK; red for fault; amber for standby. But even this system was not unambiguous. Some designers

insisted that the turn-on of high-voltage should be indicated by a red light to denote that equipment has been energized, even though no fault is present. In such schemes, the potential for ambiguity is obvious.

The most recent designs take advantage of industrial-strength programmable-logic controllers (PLCs) and the almost-limitless multicolored, alphanumeric video display capabilities they possess. Far from simplifying the ergonomics, however, they make it all but mandatory that the user participate in the display design, or that a human-factors specialist be consulted to preclude a glut of uninterpretable data. PLCs now pose a real threat of converting more data into less information. It is the unambiguous and instantaneous information transfer between machine and human that is half of the goal of good ergonomic design. The other half—not surprisingly—is the optimization of the interface in the other direction: how unambiguous will the response of the machine be to a control input from the human operator? (Or, more correctly, how unambiguous will the operator's expectation be of the machine's response to a given control input?) As an example of this issue, consider for a moment the placement and shape of the accelerator and brake pedals in an automobile. These controls are so crucial to safe and effective operation that, once their shapes and positions had been established, they have never changed, while almost every other automobile control and indicator has been fiddled with in the interest of ergonomics.

4.10 Configuration control

The intention of configuration control, or configuration management as it is also called, is to make sure that a product, as it is routinely maintained or as parts are replaced following failure, remains the same as it was when it was first built. Unfortunately, in the case of a high-power transmitter, this is a virtual impossibility.

The reason for this, which will become increasingly clear as we move into specific designs, is that components, especially high-power components, simply are not being built any more in the quantity and diversity that they once were. Even low-level components, particularly special-function integrated circuits, obsolesce—or at least fall from favor. Therefore, they are no longer manufactured. (As a case in point, Motorola's entire line of high-noise-immunity gate circuits, which were used in many transmitter solid-state control systems, is no longer manufactured and has been relegated to virtual antiquity.) So, although configuration control is an admirable goal and high on many customer's and user's wish lists, for high-power transmitters the recent history of electronic component suppliers makes it just that: a wish.

How could such a situation have evolved? Can't the government, one of the biggest users of such electrical components, insist that replacement parts remain available? The answer is yes—to an extent. The problem for the high-power microwave transmitter field is that the government hasn't gone far enough. The system put in place to ensure that critical components would remain available, the joint Army-Navy (JAN) and military (MIL) standard categories of parts, has not identified enough components to populate even the simplest of transmitters. Therefore, transmitters tend to have a preponderance of non-standard parts, from a MIL-specification point of view. This means that the purchase of such parts

must be justified through a Non-Standard Parts Approval Request and approved by the government-agency customer.

Compounding the overall problem is the fact that even though there may be perfectly useful standard parts in the JAN and MIL category that are generic in nature and ideally available from more than one supplier, there may not be, at any given time, a vendor qualified to produce them. There might not even be a vendor interested in manufacturing them!

5. And Now a Word From the Sponsor

Make sure you really understand the specifications and the acceptance criteria.

Untold amounts of money and time have been squandered, and competently staffed companies have actually gone out of business, because of someone's failure to adequately heed the foregoing enjoiner.

A technical specification is written by the customer. It is a legal document. A technical proposal is written by the vendor. It can become a legal document if, at the customer's discretion, its wording should supersede that of the original specification. Nevertheless, many proposals are written in such a way that they are virtual time bombs or, in a more modern context, computer viruses. Documents of even deadlier potential are acceptance criteria, the list of tests defined by the customer that must be successfully performed before the product can be accepted and its vendor paid. The source of many of the misunderstandings created by these documents can be summed up in just three words: redundancy, ambiguity, and jargon—not enough of the first and too much of the second and third (which are often closely related).

Redundancy in documentation, as well as fault-tolerant electronic systems (transmitters included), is almost always a good thing. In fact, there is well-known rule in verbal communication that states you must say everything three times: tell the listener what you are going to say; say it; and then tell the listener what you just said. The written form of communication differs from the spoken in that a reader can read something as many times as necessary, though a listener may only hear it once. Redundancy in a written document is most effective if the same information is conveyed with different phraseology in different parts of the text, *but with the same meaning each time*. (Never underestimate the interpretive capability of the reader!)

The overarching villain in the failed specification or proposal is Ambiguity. Generally, a proposal is thought to have failed if it fails to win the job. Far more devastating, however, is the ambiguously worded proposal that wins the job by promising something the vendor never intended to deliver. A fine example of an ambiguously worded technical requirement is the following specification:

The transmitter shall be capable of operating at any frequency between 400 MHz and 500 MHz.

It looks innocuous and informative enough, doesn't it? But it isn't. Why? Because the successful bidder built a transmitter that worked at *only* 435 MHz,

which is certainly in the specified frequency range. However, what the user *really* wanted was one that would work at *all* frequencies between 400 and 500 MHz. But that was not what he literally asked for. The specification should have read: "The transmitter shall be capable of working at *all* frequencies between 400 MHz and 500 MHz."

Redundancy could have ameliorated, if not entirely eliminated, the ambiguity. If, for instance, the acceptance criteria called for testing the transmitter with input frequencies of 400, 425, 450, 475, and 500 MHz, the vendor would at least have had a clue that the earlier requirement had been poorly stated.

Another classic, but more subtle, ambiguity often encountered in technical requirements takes this form:

*The hum, noise, and distortion products shall be greater
than 60 dB below the carrier.*

Although at first glance this specification may seem common and direct enough, it actually belongs in the Ambiguity Hall of Fame because it contains several of them. First, does it mean that each of the components—hum, noise, and distortion—must meet the specification individually or when combined (in some unspecified fashion)? Second, because hum, noise, and distortion are commonly assumed to be bad things, a sophisticated reader would assume that what was wanted was for these to constitute a signal level that is smaller than one that is 60 dB below the carrier. In other words, "greater" should apply to the number 60. But look again. That is not what has been asked for. Literally, what has been asked for is that hum, noise, and distortion—in some combination or individually—should constitute signal levels that are greater than one that is 60 dB below the carrier. Maybe they should be 50 dB below the carrier, or 25 dB—anything greater than a level 60 dB below the carrier. If confronted with this ambiguity, the engineer/author might well say, "But everyone who reads it will know what I really wanted." More often than not, this will be true. But this can only make matters worse.

Now we enter into the Kingdom of Jargon. A jargon word is a common word that has been given a special meaning in certain situations or a word that is unique to a particular field or activity. In this case, the author not only took unwarranted advantage of special-purpose words—hum, noise, distortion, and decibel—but entire expressions by assuming the reader would understand the specification's intent by understanding the jargon words. Unfortunately, in this case the jargon words were used in a sentence whose structural logic was not without ambiguity. This sentence would completely lose a reader who was ignorant in electronics jargon. Readers who did know the lingo, however, might well jump to the wrong conclusion and yet be absolutely sure they understood the requirement precisely because it was written in the insider's language, which they thought they knew. A specification or proposal is no place for jargon.

An especially egregious example of jargon at work is the use, or more correctly the misuse, of the word "triaxial." This is a genuine word. Webster

defines it as "having or involving three axes." For high-voltage, high-current interconnections, high-power transmitters commonly use specially built coaxial cables that have a single center conductor that carries the high voltage and two concentric outer conductors that are insulated from each other. The inner conductor is the insulated low-voltage return path for the high-voltage circuit, and the outer conductor is customarily connected to the grounded equipment enclosures at both ends in order to extend the overall Faraday enclosure. Because there are three concentric conductors, this type of cable is all often called a "triaxial" cable, although it has but one axis that is shared by all three conductors. That is why this cable should correctly be called "coaxial," because the prefix "co-" really means "with," "together," or "joint," not "two" or "dual."

To most engineers, this argument would sound like quibbling. "Everybody knows what I really mean," they would say. But does everybody? If engineers write in a specification that they want equipment interconnected by means of a "triaxial" high-voltage cable, how can they possibly be sure that no one will misinterpret their requirement when what they want is *literally* different from what they asked for? Redundancy will help. Or so will a drawing. If engineers draw a picture of the cross section of the cable, they are more likely to get what they want regardless of what they called it.

Writing a specification or proposal is difficult, but so is reading and interpreting one. Before passing a specifications or proposals on to the reader, authors should check them again to make sure that what they actually wrote was what they actually meant. After reading specifications or proposals, readers should check them again to make sure they didn't unwittingly leap to a false conclusion.

6. What Makes Resistors Hot?

As we all know, power dissipation makes resistors hot. How hot they get is a subject to be discussed later. Why they get hot is what this chapter's all about.

Let's start with a problem. Consider Fig. 6-1. In it we see a condensed version of part of the Fig. 1-1 block diagram. It shows only the high-voltage power supply, the energy-storage capacitor bank, a load (in the form of a diode-gun microwave-amplifier tube), a pulse modulator (in the form of a single-pole, single-throw switch), and two resistors, one between the power supply and capacitor bank and the other between the capacitor bank and load.

The current through the load is pulse-modulated by opening and closing the switch. When the switch is closed for interval T_{on} , the current, I_2 , which is limited by the diode load, is I_{peak} . This is also the current flowing through the resistor R_2 . During the interpulse interval, shown as T_{off} , the switch is open. The current during this interval is zero. The pulses recur at a constant rate called the pulse-repetition frequency (PRF). The interval between the leading edges of successive pulses is called the pulse-repetition interval (PRI). The duty factor of this waveform is the ratio of the "on" interval to the PRI. Because our capacitor C_{huge} is arbitrarily large, and because our waveform has presumably been going since the beginning of time, it is safe to assume that all of the current delivered to the load during the pulse came from the storage capacitor. It is also safe to assume that the current from the power supply is continuous, non-variant, and equal to the time-averaged value of the pulse current, which is the product of I_{peak} times the duty factor. These conditions, and the fact that there can be no average current

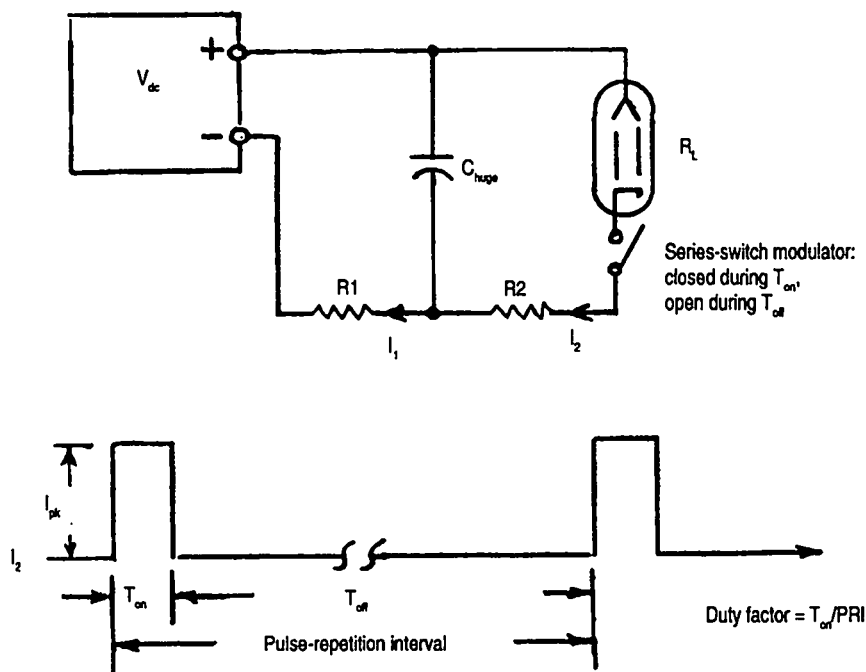


Figure 6-1. Which resistor gets hotter?




Figure 6-2. Recurrent rectangular waveform.

through a capacitor in any case, mean that the average current is the same in $R1$ and $R2$. Now, here's the problem: if they both have the same resistance, which one gets hotter? (You might be amazed at how many engineers do not know the correct answer, even if you do.)

The fact that the average current through both resistors is the same does not mean that the power dissipation is the same, because average current is not what causes power-dissipation in resistors. (If it did, pure alternating current would cause no dissipation, because its average value is zero.) Average current is what causes power dissipation in unidirectional, non-linear devices like semiconductor diodes and ionized plasma. These devices tend to have a voltage drop that, to a first approximation, is independent of current. What causes power dissipation in resistors is root-mean-square (RMS) current. Unlike any other circuit element, the voltage across a resistor is proportional to and in phase with the current through it. Power dissipation, then, will be proportional to the square of the current. Therefore, $R2$ will always get hotter than $R1$ —unless the duty factor reaches unity, at which point the current flowing through both resistors will be unidirectional and non-variant, which is the only wave shape that has identical RMS and average values. The term dc, by the way, is not adequate to define a non-variant unidirectional current, because the wave shape in Fig. 6-1 is dc for all duty factors. The current still only goes in one direction, never reversing. (Do not be fooled by the Fourier series, which correctly characterizes the rectangular-pulse periodic waveform as an average term and an infinite series of sinusoidal ac terms that are harmonically related to the PRF. They all exist simultaneously, and their sum is unidirectional.)

Figure 6-2 shows a recurrent rectangular pulse waveform. The time-averaged and RMS values for the current can be found by using these simple equations:

$$I_{avg} = \frac{\int_0^{T_p} I_{pk} dt}{T_i} = \frac{I_{pk} T_p}{T_i} = I_{pk} DF$$




Figure 6-3. Recurrent rectified sine wave.

$$I_{RMS} = \sqrt{\frac{\int_0^{T_p} I_{pk}^2 dt}{T_i}} = \sqrt{I_{pk}^2 \frac{T_p}{T_i}} = I_{pk} \sqrt{\frac{T_p}{T_i}} = I_{pk} \sqrt{DF}$$

$$\frac{I_{RMS}}{I_{avg}} = \frac{\sqrt{DF}}{DF} = \frac{1}{\sqrt{DF}}$$

$$\frac{I_{RMS}^2}{I_{avg}^2} = \frac{1}{DF}$$

The important thing to note is that the ratio of RMS and average current varies inversely as the square root of the duty factor. Even more importantly, the square of this ratio, which is directly proportional to heating, varies as the inverse of the duty factor. This means in our example of Fig. 6-1 that if the duty factor had been 0.001 (or 0.1%)—not an uncommon duty factor for simple pulse radar systems—resistor R_2 would have to dissipate 1000 times as much power as R_1 .

Figures 6-3, 6-4, and 6-5 show the derivations for the corresponding values for the recurrent rectified sine wave, the recurrent triangular waveform, and the recurrent exponential decay waveforms, respectively.

Refer to Fig. 6-3 for the following derivation of RMS and average current values:

$$I_{avg} = \frac{T_p}{\pi T_i} \int_0^{T_p} I_{pk} \sin \frac{\pi t}{T_p} dt = \frac{T_p I_{pk}}{T_i \pi} \left[-\cos \frac{\pi t}{T_p} \right]_0^{T_p}$$

$$= \frac{T_p I_{pk}}{T_i \pi} [-\cos \pi - (-\cos 0)] = \frac{T_p I_{pk}}{T_i \pi} [-(-1) - (-1)] = \frac{2 T_p I_{pk}}{\pi T_i}$$

$$= \frac{2}{\pi} I_{pk} DF$$




Figure 6-4. Recurrent triangular waveform.

$$\begin{aligned}
 I_{RMS} &= \sqrt{I_{pk}^2 \frac{T_p}{\pi T_i} \int_0^{T_p} \sin^2 \frac{\pi t}{T_p} dt} = I_{pk} \sqrt{\frac{T_p}{\pi T_i} \frac{1}{2} \left[\frac{\pi t}{T_p} - \sin \frac{\pi t}{T_p} \cos \frac{\pi t}{T_p} \right]} \\
 &= I_{pk} \sqrt{\frac{T_p}{T_i} \frac{\pi}{2\pi}} = \frac{I_{pk}}{\sqrt{2}} \sqrt{\frac{T_p}{T_i}} \\
 &= \frac{I_{pk}}{\sqrt{2}} \sqrt{DF}
 \end{aligned}$$

$$\frac{I_{RMS}}{I_{avg}} = \frac{\pi}{2\sqrt{2}\sqrt{DF}} = \frac{1.11}{\sqrt{DF}}$$

$$\frac{I_{RMS}^2}{I_{avg}^2} = \frac{1.23}{DF}$$

Note that at unity duty factor the average value is the familiar $2/\pi$ times the peak value, and the RMS value is the even more familiar peak value divided by the square root of 2. Even at unity duty factor, however, the ratio of RMS to average current is greater than unity because the rectified sine wave is not a constant value during its "on" interval.

Figure 6-4 shows the recurrent triangular wave, which is even farther away from a constant value during the "on" time than the rectified sine wave. The triangular wave manifests an even greater ratio of RMS-to-average current values. The derivation of these values is as follows:

$$\begin{aligned}
 I_{avg} &= \frac{1}{T_i} \int_0^{T_p} \frac{I_{pk}}{T_p} t dt = \frac{I_{pk}}{T_i} \left[\frac{I_{pk} t^2}{2t_p} \right]_0^{T_p} = I_{pk} \frac{T_p^2}{2T_i T_p} = \frac{I_{pk} T_p}{2T_i} \\
 &= \frac{I_{pk}}{2} DF
 \end{aligned}$$




Figure 6-5. Recurrent exponential waveform.

$$I_{RMS} = \sqrt{\frac{1}{T_i} I_{pk}^2 \int_0^{T_p} \frac{t^2}{T_p^2} dt} = \sqrt{\frac{1}{T_i} I_{pk}^2 \left[\frac{t^3}{3T_p^2} \right]_0^{T_p}} = I_{pk} \sqrt{\frac{T_p^3}{3T_i T_p^2}} = \frac{I_{pk}}{\sqrt{3}} \sqrt{\frac{T_p}{T_i}}$$

$$= \frac{I_{pk}}{\sqrt{3}} \sqrt{DF}$$

$$\frac{I_{RMS}}{I_{avg}} = \frac{2}{\sqrt{3} \sqrt{DF}} = \frac{1.555}{\sqrt{DF}}$$

$$\frac{I_{RMS}^2}{I_{avg}^2} = \frac{1.33}{DF}$$

The waveform of Fig. 6-5, the recurrent exponentially decaying waveform, is also quite common in transmitter-type circuits. It has no "on" time as such. The decay time-constant plays a similar role. The RMS-to-average current ratios shown are true as noted only for repetition intervals that are very large with respect to the decay time-constant. The derivation of their values is as follows:

$$I_{avg} = \frac{1}{T_i} \int_0^{T_i} I_{pk} e^{-t/\tau} dt = \frac{\tau}{T_i} \left[I_{pk} e^{-t/\tau} \right]_0^{T_i} = I_{pk} \frac{\tau}{T_i}$$

$$I_{RMS} = \sqrt{\frac{1}{T_i} \int_0^{T_i} I_{pk}^2 e^{-2t/\tau} dt} = \sqrt{\frac{\tau}{2T_i} I_{pk}^2} = I_{pk} \sqrt{\frac{\tau}{2T_i}}$$

$$\frac{I_{RMS}}{I_{avg}} = \sqrt{\frac{T_i}{2\tau}}$$




Figure 6-6. RMS values for short-circuit currents in ac circuits.

$$\frac{I_{RMS}^2}{I_{avg}^2} = \frac{T_i}{2\tau}$$

The waveforms of Fig. 6-6 are especially interesting for a couple of reasons. Unlike the others, they are ultimately ac waveforms. They result from short-circuiting an ac source, which is exactly what unintentional flash arcs in the high-power microwave tube or intentional breakdowns of an electronic crowbar switch will do to the ac primary power of the high-power transmitter shown in Fig. 1-1.

Where the RMS value of the ac component can be given as

$$\frac{I_{pk}}{\sqrt{2}}$$

and where the RMS value of maximum offset is $I_{pk'}$ then the initial RMS value of asymmetrical current is

$$I_{RMS} = I_{pk} \sqrt{I^2 + \left(\frac{1}{\sqrt{2}}\right)^2} = I_{pk} \sqrt{\frac{3}{2}}$$

The magnitude of the short-circuit current is limited by the reactance of the source inductance, which is derived predominantly from the leakage reactance of the step-up and/or step-down transformers that precede the high-voltage rectifier assembly. Even though there will always be series resistance as well, when it is combined as the square root of the sum of the squares with the inductive reactance, it invariably has a negligible contribution to the total source impedance. (The resistive component *will* have an effect on the nature of the current waveform, however, as we shall see.)

Have you ever watched fuses blow? If you have, you have probably wondered why sometimes the event is accompanied by a flash of light that results in the fuse's glass enclosure being completely blackened on the inside and other times, under what appear to be the same conditions, the event causes only a tiny break in the fuse link. The difference may not be due to any physical difference between the fuses. It may be due to the nature of the current flowing through it.

As mentioned before, the short-circuit current will be limited by the source inductive reactance, which means that it will also lag the source voltage by 90° yet have the same wave shape as the voltage did 90° earlier. Suppose the arc that short circuits our system occurs when the supply voltage is at an instantaneous maximum. The short-circuit current will commence with the same wave shape that the voltage had 90° earlier when it was passing through the zero-voltage point. The current wave shape will be sinusoidal and symmetrical until the fast-interrupt switch clears the line. The peak value of current will be the peak value of voltage divided by the source inductive reactance and will generally be 10 to 15 times greater than the normal full-load current. But what if the short circuit does not occur at a voltage maximum? Suppose, instead, that it occurs at a voltage minimum, when it is passing through zero. The short-circuit current will once again commence with the wave shape that the voltage had 90° earlier, which, in the case illustrated, was a negative peak. If the current has this shape starting from zero, it will continue to grow as a cosine wave until its instantaneous maximum is twice the peak value. This short-circuit current is fully asymmetrical and has a unidirectional maximum offset equal to the peak value of current. This offset component will exponentially decay with a time-constant of L/R , where L and R are the source inductance and resistance, respectively. We know what the RMS value of the symmetrical current is: the peak value divided by the square root of 2.

At the beginning of the fault current, the fully asymmetrical waveform is a composite of a constant-value dc term, which is equal to the peak value of symmetrical current and the symmetrical current itself. This condition is especially true if the L/R time-constant is long. The RMS value of this or any other composite waveform comprising components at different frequencies can be evaluated by taking the square root of the sum of the squares of the RMS values of the individual components. For this waveform the value is

$$I_{pk} \sqrt{\frac{3}{2}},$$

which is $\sqrt{3}$ times as large as the RMS value of the symmetrical waveform. This means that the initial power dissipation in a fuse, or the initial magnetic force in a circuit breaker, can differ by a factor of 3, depending on when in the voltage cycle the short circuit occurs. That's one reason why the same type of fuse can sometimes put on a show when it blows and at other times give up the ghost quietly.

7. High-Voltage Insulation

More than any other factor, what separates high-power design from low-power is insulating for high voltage. Thermal management, on the other hand, is just as important for low-power systems as for high-power ones. Poor attention to thermal design can make a low-power system just as hot as a high-power system that suffers from a comparable lack of attention. The physical attributes are also proportional, in general. What is it, then, that separates a low-power system from a high-power one in terms of insulating for high-voltage?

Figure 7-1 gives us a point of departure. It is Paschen's curve for electrical breakdown in air, the most commonly used insulating medium. It plots the voltage between electrodes that will cause electrical breakdown, or ionization of the air between them, as a function of the product of the barometric pressure and electrode spacing. The most immediately recognizable piece of information from the curve is the absolute voltage minimum, approximately 340 V, below which it is not possible to produce breakdown in air regardless of spacing or pressure. But what about electrical systems of less than 340 V where arcing has been seen? For instance, what about the arc discharges that we have all observed coming from the terminals of a car's 12-V storage battery when we need to jump-start it? The arcing does not come about when we connect the battery to an external circuit. No matter how slowly we move the circuit lead to the battery terminal, we cannot generate an arc. Once the circuit is complete, however, current and self-inductance are established in the circuit. Now any attempt to break the circuit will almost always produce an arc—but not because of the original 12 V. The voltage across an inductor, expressed quantitatively as $-L di/dt$, is qualitatively whatever voltage is required to maintain whatever current is flowing in the inductance. And the source of the voltage is the energy stored in its magnetic field, $1/2 LI^2$. This circuit-sustaining voltage can and will be more than enough to ionize air. (Indeed, this is how arc-welding machines work.)

Therefore, designers of systems whose operating voltages are less than 340 V do not have to worry about electrical insulation but only about preventing conductors or electrodes from coming into contact with each other. These systems, then, are most properly in the low-power domain. Given that most high-power systems will operate at voltages above 340 V (except for a precious few highly modular solid-state systems), Paschen's curve can also tell us just how good an electrical insulating medium air is for higher voltages. For instance, note the data point at the upper-right-hand corner of the curve. It shows that breakdown for electrodes spaced by one centimeter at one atmosphere of barometric pressure (and 25°C ambient temperature) will occur at 30,000 V. (At one-inch spacing under similar conditions, breakdown occurs at 79,000 V.) Those who have had any practical experience with high-voltage design will recognize that these are hardly the design criteria most commonly used. (The far more common guideline is one inch of separation for every 10,000 V.) Does the curve lie? No it does not. But it is based on conditions that are almost never realized in practice: that the conductors are an infinite plane, parallel, and perfectly smooth. When these




Figure 7-1. Paschen's curve for electrical breakdown of air.

electrode conditions are approached, air will perform as advertised.

Breakdown voltage is also influenced by atmospheric pressure. For two parallel-plane electrodes whose spacing is uniform, breakdown voltage decreases as barometric pressure decreases until the 340-V minimum value is reached. This phenomenon is due to the fact that as the molecular density of the air is reduced, there is greater likelihood that a free ion can traverse the space between electrodes without running into something. However, as pressure is reduced even further, the required voltage for breakdown increases once again. This is because a more limited number of air molecules make ionization more difficult. Note in the inset graph in Fig. 7-1 that as altitude increases up to 50,000 ft, where atmospheric pressure is 87 mm of mercury, voltage hold-off performance steadily deteriorates, and the deterioration is more severe for large spacings than for small.

Far more common than parallel-plane electrodes, which resemble the interior of a smooth-sided box, are combinations with different geometries, where some electrodes resemble spheres or points (or spherical or pointed ends of cylindrically shaped electrodes), and others resemble infinite planes. Figure 7-2 illustrates what happens to electric-field intensity in such geometries. Between parallel planes, the electric fields are uniform everywhere in the space between them and are normal to the surfaces. Equipotential lines (in the two-dimensional representation) or planes (in the three-dimensional representation) are equally




Figure 7-2. Determination of maximum voltage gradient between spherical and planar electrodes.

spaced and are normal to the electric-field lines, which means they are parallel to the electrodes.

However, when one of the electrodes is spherical in shape (circular in the two-dimensional representation), note that the equipotential lines on the right-hand side of the figure come closer together the nearer they are to the surface of the sphere. The voltage difference between any two equipotential lines is the same, which means that the electric-field gradient or strength increases as one nears the surface of the sphere or point and can be many times as great as the field strength at the surface of the planar electrode. The exact expression for peak voltage gradient electric field of this configuration can be stated as

Table 7-1. Voltage gradients between spherical and planar electrodes.

R = constant		D = constant		
d	Peak electric field	R	d = D-R	Peak electric field
0.01R	1.007 V/d	0.01D	0.99D	100 V/D = V/R
0.1R	1.07 V/d	0.1D	0.9D	10.6 V/D
0.5R	1.37 V/d	0.1D	0.9D	10.6 V/D
R	1.78 V/d	0.3D	0.7D	4.3 V/D
2R	2.7 V/d	0.4D	0.6D	3.7 V/D
5R	5.6 V/d = 1.1 V/R	0.5D	0.5D	3.6 V/D
10R	1.5 V/R	0.6D	0.4D	3.75 V/D
20R	1.03 V/R	0.7D	0.3D	4.4 V/D
100R	1.005 V/R	0.8D	0.2D	5.9 V/D
Inf.	V/R	0.9D	0.1D	10.8 V/D

$$\text{Gradient} = \frac{\left\{ \frac{2d}{R} + 1 + \left[\left(\frac{2d}{R} + 1 \right)^2 + 8 \right]^{1/2} \right\} V}{4d},$$

where d is the distance between electrodes, R is the radius of the spherical electrode, and V is voltage (see Fig. 7-2). As d/R approaches infinity,

$$\text{Gradient} \Rightarrow \frac{\left(\frac{2d}{R} + \frac{2d}{R} \right)}{4d} V = \frac{4d}{4dR} V = \frac{V}{R}.$$

The approximate expression for peak voltage gradient is

$$\text{Gradient} = 0.9V \frac{R+d}{Rd}.$$

Table 7-1 shows how the electric-field intensity at the surface of the sphere varies as the spacing between the sphere and the plane is varied. When the sphere is close enough to the plane so that the spacing is small compared to the radius of curvature, both electrodes appear to be relatively planar (much as the surface of the Earth appears flat to those of us standing on it, even though many of us believe the Earth to be spherical). When this is the case, the electric-field intensity is very near to V/d , where d is the spacing.

However, as the spacing is increased, the roundness of the sphere becomes increasingly apparent. When the spacing is large with respect to the radius of curvature, the gradient at the surface of the sphere approaches V/R , where R is the radius of curvature, and has nothing to do with the spacing. Thus, if V/R for a given conductor exceeds 79 kV/in. (or 30 kV/cm), the air at its surface will ionize, producing what is called a corona, or partial discharge. If the spherical electrode is already a long way from the planelike conductor, moving it farther away will not increase the corona-inception voltage. In most cases, the corona is self-limiting because the ionized air is an electrical conductor that, in effect, increases the radius of curvature of the electrode to the point where the electric-field gradient at its outer surface is no longer high enough to ionize any more air.

Figure 7-3 shows the same information in a slightly different way. It shows that the electric-field gradient, or strength, normalized either to V/d or V/R . The gradient reaches a maximum of 1.78 when the spacing and radius are the same and asymptotically approaches V/d and V/R as the normalized dimensions R/d and d/R increase from unity. What is interesting is how rapidly the asymptotes are approached. When either normalized dimension is greater than 5, it is almost all the way there.

Situations are occasionally encountered when electrode geometry is better described as sphere-opposing-sphere, or cylinder-opposing-cylinder (such as parallel-wire transmission line), as shown in Fig. 7-4. When the two electrodes are close together, they behave as though they were both planar, with the electric-field gradient approaching V/d as shown before. When they are separated by an




Figure 7-3. Electric-field intensity at surface of spherical conductor opposing plane.

amount that is large with respect to the radii of curvatures, electric-field strength increases at the surface of both conductors. What is surprising to many is the fact that the voltage between conductors required for surface corona is twice as great for sphere-to-sphere as for sphere-to-plane geometry because the aggravated condition at the spherical surfaces is now, in effect, divided in half. (The corona will actually start first at the surface of the conductor that is negative with respect to the other.) Notice the symmetry and the shape of the equipotential lines. At the mid-point between conductors a conducting plane could be inserted without affecting the situation at all. For electrodes in this configuration, the exact expression for peak voltage gradient can be stated as

$$\text{Gradient} = \frac{\left\{ \frac{d}{R} + 1 + \left[\left(\frac{d}{R} + 1 \right)^2 + 8 \right]^{1/2} \right\} V}{4d} .$$

As d/R approaches infinity,

$$\text{Gradient} \Rightarrow \frac{\left(\frac{d}{R} + \frac{d}{R} \right)}{4d} V = \frac{2d}{4dR} V = \frac{V}{2R} .$$

The approximate expression for peak voltage gradient for two spherical electrodes is

$$\text{Gradient} = 0.9V \frac{R + \frac{d}{2}}{Rd} .$$




Figure 7-4. Determination of maximum voltage gradient between two spherical electrodes.

Table 7-2 shows how the electric-field intensity at the surface of the spheres varies as the spacing between the spheres changes.

Figure 7-5 shows plots of actual breakdown performance of spherical electrodes whose diameters increase from 6.25 to 25 cm and whose spacing increases from 0 to 40 cm. Notice that they all behave similarly when they are 1 cm apart; they all break down at about 22 kV. In other words, their average-voltage gradient is 22 kV/cm (which isn't 30 kV/cm, but it's close). All electrodes, including even the 2-cm-diameter ones, behave almost as if they were infinite planes. However, as the electrodes are separated by increasing distance, the average-voltage gradient at which breakdown occurs diminishes as the curves tilt toward the horizontal, or $V/2R$ -asymptote, direction.

Figure 7-6 shows the improvement gained by replacing air with transformer oil. (Sometimes dielectric gases like nitrous oxide and sulfur hexafluoride are used, but in applications where weight and mobility are not important, oil is usually preferred.) The trends related to electrode geometry are similar for oil and air, except that for oil, the V/d asymptote is at 200 kV/cm instead of 22 kV/cm for air—almost 10 times the dielectric strength. As the spacing-to-radius ratio is increased, the tilting to the horizontal is, if anything, more abrupt than it is for air. Note the effect of a needle-pointed gap. Because the radius of curvature is so

Table 7-2. Voltage gradients between two spherical electrodes.

R = constant		D = constant		
d	Peak electric field	R	d = D-R	Peak electric field
0.01R	1.003 V/d	0.1D	0.8D	5.8 V/D
0.1R	1.073V/d	0.15D	0.7D	4.3 V/D
0.5R	1.2 V/d	0.2D	0.6D	3.7 V/D
R	1.4 V/d	0.25D	0.5D	3.6 V/D
2R	1.8 V/d	0.3D	0.4D	3.75 V/D
5R	3.2 V/d	0.4D	0.2D	5.9 V/D
10R	5.6 V/d = 0.56 V/R			
20R	10.5 V/d = 0.53 V/R			
100R	50.5 V/d = 0.505 V/R			
Inf.	0.5 V/R			




Figure 7-5. Breakdown performance of spherical electrodes in air ($t = 25^{\circ}\text{C}$, $p = 1 \text{ atm}$)

small to begin with, the characteristic has no V/d component to speak of. Therefore, it is horizontal-leaning over the entire range of gap spacings. The needle-gap geometry is almost the least desirable of all. (The absolute least desirable is the geometry where the needle opposes the plane. Unfortunately, this geometry is the one approximated by unrelieved screw bodies and many forms of electrical terminals. In such cases, the screws must either be smoothed off or recessed.)

In order for it to be an effective insulator, transformer oil must be dry, which means low in water content. A standard test for oil dryness is the American Society for Testing and Materials (ASTM) voltage-breakdown test, which is illustrated in the upper left of Fig. 7-6. The test apparatus uses cylindrical, flat-faced electrodes spaced 0.1 in. apart. For oil dryness to be minimally acceptable, breakdown must not occur with 30 kVac (60 Hz) applied between electrodes. This is the so-called 30-kV oil figure of merit.

Oil, like most gases and vacuum, has higher dielectric strength for short pulses than for dc or 60-Hz ac. When tested with a standard lightninglike short-duration waveform (1.5- μs rise from zero to peak, 40- μs fall from peak to half-amplitude), oil will withstand 2 to 3 times the dc or 60-Hz ac breakdown voltage. This time-dependency has been ascribed to the inertia of conducting dipoles within the oil. They must be aligned end-to-end before the equivalent of ionization occurs.

Solid dielectrics are used in many insulating applications, especially where they also play a mechanical role as well. Such dielectrics are the ceramic or glass insulators used to separate high-power vacuum-tube electrodes that have large potential differences. In these cases, the insulator also functions as an extension of the vacuum-tight envelope of the tube. High-voltage transmission cables use plastic dielectrics such as polyethylene or Teflon.

When dealing with solid dielectrics, special care must be taken to assure that there is no trapped air between the metallic conductors and the solid insulation. Semiconducting layers are often used in such applications to act as a transition between conductor and insulator in an effort to short-circuit any trapped air. The terminations at either end of such cables are especially prone to the problems of trapped air. Often their castings are made of an epoxy resin that has been embedded with metallic shapes that linearize the electric-field gradient along the surface of the insulator. However, unless these castings are made under the best of conditions—usually in a vacuum—they too can have trapped air bubbles.

Figure 7-7 illustrates the plight of an air-filled void within a solid dielectric. The problem is double-barreled. First, the solid dielectric has a higher relative permittivity (dielectric constant) than air, ranging from 2 to 3 for most plastics and up to 9 for high-alumina ceramic. Therefore, an air-filled void in such material will function as a capacitive voltage divider. As such, the voltage stress




Figure 7-6. Insulation performance of dry transformer oil at $t = 25^{\circ}\text{C}$.




Figure 7-7. Effect of trapped-air void in solid dielectric insulator.

across the void, per unit of length, will be the product of the solid dielectric electric-field strength (V/d) and its relative permittivity (ϵ_R). This is bad enough, but the second problem is the electric field. It will be higher in the solid dielectric to begin with because of its greater dielectric strength (which was one of the reasons for using it in the first place). It is not unreasonable to expect that the air in the void will be ionized at the intended working voltage. The ionized air will be superheated and will wish to expand, a desire that it usually satisfies by blowing the casting apart.

In general, designing for low corona at high voltages requires that electric fields be made as uniform as possible, which means a designer needs to approach the condition of parallel-plate geometry as closely as possible. When this is physically impractical, it is mandatory to remember the criterion that the operating voltage, V , divided by the radius of curvature of the smallest-diameter electrode in a system, R , must be less than the breakdown-voltage gradient of air (79 kV/in., or 30 kV/cm, at sea level and room temperature). And this is for electrodes spaced by many times the radius of curvature. As spacings are decreased, the operating voltage must also be decreased by a factor that reaches a maximum of 1.78 when spacing is equal to radius of curvature. For spacings smaller than this, the parallel-plate criterion of V/d becomes increasingly dominant and flashover occurs at the same voltage at which corona would begin.

One of the more difficult problems in high-voltage engineering is the termination of a high-voltage coaxial cable, as shown in Fig. 7-8. The problem is not how to terminate the cable inner conductor but how to terminate the shield. High-performance, high-voltage coaxial cables use highly flexible rubberlike dielectrics with carbon-filled, semiconducting layers that are located between the outer sur-

face of the copper inner conductor and the inner diameter of the dielectric and between the outer diameter of the dielectric and the outer shield. This is done to make sure that there is no trapped air in these interfaces.

Why such concern? Consider a typical high-voltage cable that has a 1-in. shield diameter and whose inner conductor has a diameter of about 0.2 in. The spacing, d , between inner and outer conductors is 0.4 in., and the radius of the inner conductor, r , is 0.1 in. The cable is rated for over 100 kV. If the conductors were parallel plates, the electric-field intensity in the dielectric would be 100 kV/0.4 in., or 250 kV/in., just for reference. In coaxial geometry, however, the electric field is not constant in the radial direction. Its strength is inversely proportional to the radial distance measured from the center line. The total voltage across the cable is given as

$$V = \int_{R=r}^{R=r+d} \frac{k}{R} dR,$$

where k is a constant. The integration yields

$$V = k \ln(r+d) - \ln(r) = k \ln\left[\frac{(r+d)}{r}\right] = k \ln\left(1 + \frac{d}{r}\right).$$




Figure 7-8. The problem of terminating a high-voltage coaxial cable.

The value for the constant k can be derived as follows:

$$k = \frac{V}{\ln\left(1 + \frac{d}{r}\right)} = \frac{100 \text{ kV}}{\ln\left(1 + \frac{0.4}{0.1}\right)} = \frac{100 \text{ kV}}{\ln(5)} = 62 \text{ kV}.$$

At the surface of the inner conductor, the electric-field intensity is $k/r = 62 \text{ kV}/0.1 \text{ in.}$, or 620 kV/in. At the outer diameter of the dielectric, the electric-field intensity is $k/(r + d) = 62 \text{ kV}/0.5 \text{ in.}$, or 124 kV/in. But this is the field strength in the solid dielectric, which has a dielectric constant of about 3. A pocket of air trapped between dielectric and shield would see a field strength of $3 \times 124 \text{ kV/in.}$, or 372 kV/in. , and air will break down when the field strength exceeds 79 kV/in.

When terminating the shield as shown in the top half of Fig. 7-8, the equivalent of an air pocket is created right at the point where the shield stops but the solid dielectric continues. For the cable discussed above, the air in the immediate vicinity will ionize, producing localized corona that may or may not lead to breakdown to the terminated center conductor. In any case, it will do the dielectric no good and will eventually lead to cable failure at the circumferential stress point.

Successful coaxial cable terminations more closely resemble the cross-section shown in the lower half of Fig. 7-8. In this case, a metal stress cone is slipped under the outer braid. The cone flares back to a disclike flange, which becomes the outer-conductor connection. This connection will not work in air because of the mismatch of dielectric constants and dielectric strengths. If, however, the cable end can be potted by means of a void-free epoxy casting or room-temperature-vulcanizing (RTV) silicone that more closely matches the cable dielectric in permittivity and strength, the transition can be made with uniform electric-field strength along the outer surface of the casting. In such a configuration, the field strength is safely below the breakdown strength of air. (This type of transition can also work if immersed in transformer oil or if the casting is hollow and filled with oil.)

8. Pulse Modulators

For many designers, the most technically challenging subsystem of any high-power transmitter is its pulse modulator. Even though not all transmitters require one, the complete transmitter engineer can't know too much about them.

To begin, there are some system-level considerations that govern which of the two basic types of pulse modulators are appropriate for a given high-power microwave tube application: cathode pulsers and control-electrode pulsers. Of course, in the latter case it is necessary that the microwave tube have a control electrode in its electron gun. Such control electrodes include

- a modulating anode (sometimes called an isolated anode),
- a focus electrode, and
- a control grid.

A brief description of each of the major types of pulse modulators follows in this chapter, along with a discussion of the various considerations to keep in mind about electron guns (e-guns).

8.1 Cathode-controlled electron gun

Figure 8-1 shows the cross section of a diode (or cathode-controlled electron gun), which is the simplest of all electron guns. It consists of a thermionically emitting cathode, a heater, a focus electrode, and an anode. The heater raises the cathode temperature high enough for it to become an efficient emitter of electrons. The focus electrode is electrically connected to the cathode, and its shape




Figure 8-1. Cross-section view of a diode, or cathode-controlled electron gun.




Figure 8-2. High-power microwave-tube electron gun with modulating anode.

is important to the electrostatic lens formed between cathode and anode. (It is also the terminus of most of the high-voltage arcs, reducing the likelihood of arc damage to the cathode surface itself.) The anode is usually the entrance to the beam tunnel of the amplifier device. There is only one electrode to which pulse modulation can be applied and that is the cathode. (The cathode heater power source must, of course, be capable of being floated atop the cathode modulation pulse.)

8.2 Modulating anode

Some time ago, it occurred to engineers at EIMAC, who designed and built large CW klystrons for UHF television service and tropospheric-scatter communications, that if the electrode serving the function of the anode in the diode gun were insulated from both the cathode and the body of the tube, some interesting properties might be obtained. The original intent of their "isolated anode" was to mitigate the effects of internal gun arcing. If, for instance, the anode was maintained at a voltage near ground by connecting it to ground through a resistor, an arc between cathode and the isolated anode would simply pull the anode down to cathode potential and result in nothing more than the shutting off the beam current. If the electrode can be used to shut off beam current, it can also be used to turn it on, hence the alternative appellation "modulating anode."

The modulating anode is an electrode of cylindrical geometry placed between the cathode and drift-tube entrance. The example shown in cross-sectional view




Figure 8-3. Cross-section view of electron gun with insulated focus electrode.




Figure 8-4. Cross-section view of electron gun with intercepting grid.

in Fig. 8-2 is of an electron gun whose cathode operates at approximately 100 kV negative with respect to body (ground). The pulse voltage applied to a modulating anode is usually a large fraction of the cathode voltage (50% or more) and is often pulsed to within a few hundred volts of ground during the beam-current pulse. The particular geometry illustrated is typically operated at 70% cathode voltage (70-kV pulse with respect to cathode, 30-kV negative with respect to ground). The convergent electron beam, the outer edges of which have been shown with dashed lines, is magnetically and electrostatically focused so that only a tiny fraction (0.1%) of the total beam current is intercepted by the modulating anode. Therefore, it takes virtually no power to drive it. The load the anode presents to the pulsed-voltage source is predominantly capacitance, made up of its internal capacitance and the self-capacitance of the modulator itself (which, as we will see, can be the largest component).

8.3 Insulated-focus electrode

As shown in Fig. 8-1, the focus electrode in a diode-type electron gun is connected to the cathode. As shown in Fig. 8-3, it can also be insulated from the cathode, in which case a voltage can be applied between the focus electrode and cathode. During the interpulse interval, this voltage can be sufficiently negative to inhibit the flow of electrons leaving the cathode, and during the pulse it can either be at cathode voltage or even positive with respect to the cathode. As with the modulating-anode electron gun, the cathode voltage in this device is continuously applied. The voltage swing required between beam cutoff and full-beam conduction is less for the focus electrode than for the modulating anode, but the effect on beam shape during the rise-and-fall intervals is greater.

8.4 Intercepting control grid

Requiring even less voltage swing between beam cutoff and full conduction is the intercepting control grid, illustrated in Fig. 8-4. The grid is a radially connected mesh of concentric circles that is set directly in the path of the electron




Figure 8-5. Cross-section view of electron gun with control and shadow grids.




Figure 8-6. Cut-away view of electron gun with bonded-grid control electrode.

beam. Therefore, it is going to intercept some fraction of the total cathode current whenever its voltage is positive with respect to the cathode. This interception current not only subtracts from the available interaction beam current, but it causes the grid wires to heat and, eventually, to become secondary electron emitters, which severely limits the pulse energy and average-power capabilities of tubes that use them.

8.5 Shadow-gridded gun

The shadow-gridded gun, illustrated in Fig. 8-5, uses an electrode called the shadow grid that is at cathode potential, has the same pattern as the control grid, and is positioned between the control grid and the cathode so as to inhibit electrons whose paths would cause them to strike the control grid.

8.6 Bonded-grid gun

An extension of the shadow-grid concept is the bonded-grid structure, shown in Fig. 8-6. Here, an integrated assembly consisting of the grid and an insulating substrate having the same shape is directly bonded to the cathode surface. In some cases, it is even embedded in the cathode surface so that the plane of the grid is virtually flush with that of the cathode.

8.7 E-gun considerations

What, then, are the considerations that determine which type of microwave-tube electron gun is appropriate and what type of pulse modulator should go with it?

The choice is simple if the power levels involved exceed those that can be achieved with a tube that can operate with continuously applied cathode voltage. The theoretical upper limit of this voltage is approximately 150 kV. However, there are few practical systems that operate reliably at much more than 100 kV. The gray region, therefore, is between 100 kV and 150 kV. Above 150 kV there is no choice but to use a diode gun with a cathode pulser (and the shorter the pulse,

the greater the voltage that can be applied before the interval between gun arcs becomes the limitation). Between 150 kV and 200 kV, cathode pulses, whose durations can range from 40 μ s at the high-voltage end to 100 μ s at the low-voltage end, have proven practical. Between 200 kV and 300 kV, practical pulse durations shrink from 40 μ s down to 1-2 μ s. In the submicrosecond-pulse range, cathode voltages up to 1 MV are not necessarily outrageous.

On the other hand, there are situations where cathode pulsing is most inappropriate, such as for a microwave tube that is unstable if its cathode voltage is less than its intended operating voltage. (The coupled-cavity TWT is a notorious example.) Cathode pulsing of such a tube can give rise to "rabbit ears," which are periods of RF oscillation during the rise-and-fall intervals of cathode voltage as the voltage passes through regions of instability. Tubes of such power—greater than 100 kW peak power is a fuzzy criterion—almost always have electron guns with modulating anodes. Even though these modulating anodes have pulse voltages applied to them that may be less than the total beam voltage, the electrons that enter the drift region of the tube will all have the same initial velocity. This velocity is determined only by the total accelerating voltage, which is the voltage between cathode and ground. (Guns whose modulating-anode voltage is significantly smaller than the cathode voltage—say, 50% or so—are called velocity-jump guns because the electron-beam velocity jumps as it enters the drift space beyond the modulating anode.)

The tube with the modulating-anode gun and pulser is also superior to the cathode modulator in high-duty-factor, high-average-power applications because the modulator, as we will see, handles only a small fraction of the total beam power. These advantages shrink as the system pulse repetition rate increases, and at some point greater than 5000 pps—another fuzzy criterion—the microwave-tube gun must have a low-voltage control grid and a grid modulator to go with it. Worthy of consideration in the gray region between high-PRF grid control and high-average-power, high-duty-factor modulating-anode control is control by means of the focus electrode, which is located physically and electrically between them.

9. Cathode Pulsers: Line-Type Modulators

The first type of modulator to be discussed will be the cathode pulser, specifically the line-type modulator. But before we launch into it, please ponder the following elementary question:

In a simple series circuit that comprises load impedance, no reactance, and a voltage generator with internal impedance, what value of source impedance will permit maximum power delivery to the load?

Most engineers, especially electrical engineers, know the answer to this question without even thinking. Of course the answer has to be that the source impedance must be equal to the load impedance for maximum power transfer. We learned that as college freshmen. Except that the answer is wrong. It is the correct answer to a different question: what value of load resistance will result in maximum power transfer from a generator of given source impedance? However, the correct answer to the original question is zero.

Like any power company, we want to deliver the maximum energy to our load while experiencing the minimum of internal dissipation. Our goal is 100% efficiency: all generated power is delivered to the load. We want our transmitters to be efficient, but the maximum efficiency of a generator that is matched to its load is 50%. Why? Because the simplest criteria for a matched source is that the voltage that it produces into an open circuit be twice that produced into a matched load, and the current delivered to a short circuit be twice that produced into a matched load. However, there are no high-power microwave devices that meet those criteria. Microwave power amplifiers are not matched sources of RF voltage or current, nor do we want them to be. What is meant by matching a




Figure 9-1. The simplest form of a line-type modulator (R-L-C circuit).

microwave device is usually making sure that the impedances are matched to the characteristic impedances of the transmission-line systems used to couple RF power into it and out of it. With this in mind we will begin our discussion of the most basic of all pulse modulators.

9.1 The line-type pulser

The line-type pulse modulator is sometimes referred to—not quite correctly—as a “soft-tube” modulator. It, along with the practical implementation of the magnetron microwave power oscillator, is what made microwave radar a reality. Despite its relative antiquity and brutish lack of sophistication, the line-type pulser is still the first type of high-level pulse modulator that must be considered in any new transmitter design because it has one advantage not shared by any other type of high-level modulator: it stores only as much energy as it delivers to the pulsed load each pulse. It also is switched by what can be called a “half-control” switch, which is a family of switches that can be switched from the non-conducting, forward-voltage blocking state to the forward-conducting state in response to an electrical triggering impulse. But it cannot be switched back to the initial state. (The hydrogen thyratron, or “soft-tube,” is the archetype of a half-control switch. As we will see, there are many others in the family.) If a designer decides not to use such a pulse modulator, the reasons for avoiding it must be well considered and understood.

9.1.1 *The discharge circuit*

The “line” of the line-type modulator is some form of a transmission line that comprises inductance and capacitance as the primary circuit elements and resistance as the unavoidable secondary element. (A wise person once said that all there is in the electrical firmament can be described in terms of inductance, capacitance, and resistance. An even wiser person said that it can be described in terms of characteristic impedance, time-delay, and loss.) In a line-type modulator, we will charge up the capacitance of a transmission-line segment during the interpulse interval and discharge it into the load during the pulse.

Figure 9-1 shows the most rudimentary form of a transmission line: an inductor, L , in series with a capacitor, C . In series with this line is a load resistance, R_2 . Our switch, S , is a simple toggle switch, which is in series with a diode so that current through it can only be unidirectional (dc). With the capacitance fully charged—the means of charging the capacitance during the interpulse interval will be discussed later—the switch can be closed to initiate the discharge cycle.

Figure 9-2 shows the nature of the voltage and current waveforms during discharge for three significant values of load resistance. The first is $R_2 = R = 0$, which will reveal the intrinsic properties of the network itself: its characteristic impedance and its time-delay. At time 0, when the switch is closed, all of the energy is stored in the capacitor. With resistance at zero, the discharge is that of an undamped series-resonant circuit. Current will build up in the circuit with sinusoidal wave shape, as shown in Fig. 9-2a, while the capacitor voltage diminishes in cosinusoidal fashion. At a time equal to $1/4$ of the natural period, the current will reach maximum at the same instant that the capacitor voltage passes through zero. There is no resistance, so no energy is lost. At the instant of zero




Figure 9-2. Discharge characteristics of a simple R-L-C circuit.

capacitor voltage, there is no energy stored in it, so it all must be stored in the magnetic field of the inductor. This energy is equal to $1/2 LI^2$. The total energy is unchanged, so $1/2 CV^2 = 1/2 LI^2$. The characteristic impedance can be derived as follows:

$$\frac{1}{2}CV^2 = \frac{1}{2}LI^2$$

$$CV^2 = LI^2$$

$$\frac{V^2}{I^2} = \frac{L}{C}$$

$$\frac{V}{I} = \sqrt{\frac{L}{C}}$$

The ratio of V to I , even though its maxima do not occur at the same time, is $\sqrt{L/C}$, which is the characteristic impedance of the network. The equivalent of its delay-time is the quarter-period, $\pi/2 \sqrt{LC}$. The current pulse has a duration in the forward direction that is twice the equivalent time-delay, or $\pi\sqrt{LC}$, at which point it stops because the diode is in series with the switch and prevents current reversal. Voltage across the capacitor has reached V again but in the reverse direction, so the discharge cycle terminates with all of the energy back in the capacitor.

The second resistance of interest, shown in Fig. 9-2b, is the minimum value that causes complete discharge of the network without voltage reversal, the equivalent of a "matched" load. This value of resistance is said to critically damp the network and is equal to $2\sqrt{LC}$, or twice the characteristic impedance. The load voltage has the same shape as the discharge current, of course, and reaches a maximum value of

$$\frac{V}{e\sqrt{\frac{L}{C}}}$$

after a duration of \sqrt{LC} from the time of switch closure. Up to this point, the wave shape is quasi-sinusoidal. The remainder represents quasi-exponential decay and is mostly over within a time equal to $2\sqrt{LC}$. If you were to integrate the product of load current and voltage from zero to infinity, the resulting load energy per pulse is, not surprisingly, $1/2 CV^2$, as can be shown in this derivation:




Figure 9-3. Simplified discharge circuit of pulser with artificial transmission-line pulse-forming network.

$$\begin{aligned}
 \text{Load energy} = W_L &= \int_0^{\infty} i^2 R dt = \int_0^{\infty} e^{-\frac{2t}{\sqrt{LC}}} \left(\frac{V}{L} t \right)^2 \times 2 \sqrt{\frac{L}{C}} dt \\
 &= \int_0^{\infty} \frac{2V^2}{L^2} \sqrt{\frac{L}{C}} \frac{e^{-\frac{2t}{\sqrt{LC}}}}{8LC\sqrt{LC}} \left(\frac{4t^2}{LC} - \frac{4t}{\sqrt{LC}} + 2 \right) dt \\
 &= \frac{2V^2}{L^2} \sqrt{\frac{L}{C}} \times \frac{2LC\sqrt{LC}}{8} = \frac{4}{8} CV^2 = \frac{1}{2} CV^2.
 \end{aligned}$$

The third value of resistance of interest, shown in Fig. 9-2c, is one that is many times larger than the square root of $L \times C$, or the greatly overdamped case. Here, the current is in the form of a double exponential, with a rise time dictated by the L/R time constant and a decay dictated by the RC time constant, reaching a peak value at a time approximately equal to $L/R \ln R^2 C/L$ after the switch closure. The peak value of current is slightly less than V/R , or

$$i = \frac{V}{R} (e^{-t/RC} - e^{-Rt/L}).$$

(The greater R is with respect to $\sqrt{L/C}$, the closer the peak will be to V/R). Needless to say, the waveform produced by this circuit is not a useful one to

apply to a control electrode of a microwave power tube. For most applications, the more nearly rectangular the pulse is the better.

In Fig. 9-3, the simple L - C network has been broken into cascaded segments. The total inductance is made up of series-connected inductor increments, or

$$L_{network} = \sum_1^n L ,$$

and the total capacitance is made up of capacitor increments in shunt, or

$$C_{network} = \sum_1^n C .$$

As the increments become smaller and the number of segments becomes larger, the closer the network will approximate a distributed-transmission line, whose characteristic impedance will approach

$$\sqrt{\frac{Z_{short\ circuit}}{Z_{open\ circuit}}} ,$$

which is the same as $\sqrt{L/C}$, and whose time-delay will be equal to its physical length divided by the velocity of signal propagation through it.

The load for this pulser is a microwave tube having a diode electron gun that is transformer-coupled to the pulser output terminals. The switch shown is a gas-filled "soft-tube" switch. The network is recharged after each pulse by current from an as-yet-undefined charging source. The return charging current flows through the dc resistance of the pulse-transformer primary.

Figure 9-4 shows in idealized form what the voltages and currents would be if a distributed-transmission-line pulse-forming network (PFN) were discharged into a resistive load equal to the characteristic impedance of the transmission line or into a load that is matched to the source impedance. (For those who have not forgotten the earlier insistence on zero source resistance, note that the source, although "matched" to the load, still has no internal resistance or loss. It is *characteristic* impedance that is equal to the load impedance.)

The total shunt capacitance of the distributed line of Fig. 9-4, shown as the capacitance symbol in the middle of the coaxial line segment, is initially charged to V_0 . The energy stored in this capacitance is

$$W_0 = \frac{1}{2} CV_0^2 .$$

The load resistance, RL , is connected in series with the network. When the switch is closed at t_0 , current flows in the discharge loop equal to the charge voltage divided by twice the characteristic impedance, Z_0 . The load voltage is equal to the product of this current and the load resistance, which is half of the

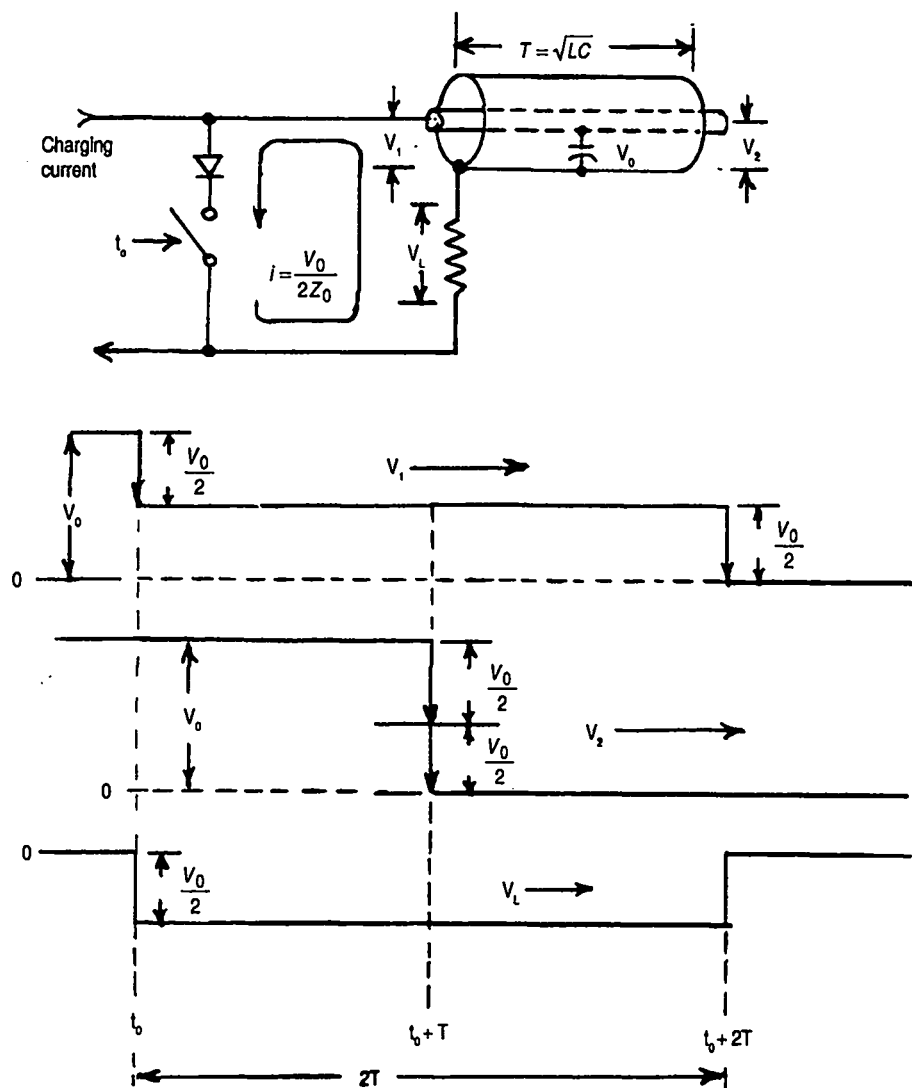


Figure 9-4. Discharge of transmission-line pulse-forming network into matched load.

network charge voltage. The sum of the voltages across the load and input increment of line impedance must be zero because there is no voltage drop across the perfect switch and diode. The load voltage is negative with respect to ground. The network voltage, at the instant of switch closure, must be positive with respect to ground by an equal amount. For this to happen, the equivalent of an incident wave equal to half of the network charge voltage begins to propagate along the transmission line to the right, toward its open-circuited far end.

In Fig. 9-4, the instantaneous voltage at the input end of the transmission-line network is shown as V_1 . This step-voltage wave results from the sequential discharge of elemental shunt capacitances delayed in time by the elemental inductances in series with them. Assuming that the dielectric between inner and outer conductors is neither liquid nor solid and that the conductors are non-ferrous, the velocity with which this wave propagates toward the open end is the speed of light, or approximately one foot per nanosecond. When this incident wave reaches the open-circuited end of the transmission line, it is reflected. The im-

pedance at the end of the line is infinite, so the reflection coefficient, Γ , which is defined as

$$\frac{R_L - Z_0}{R_L + Z_0},$$

is +1. This means that when the incident wave arrives at the end of the line, it will be reflected with the same magnitude and in the same direction as it had when it arrived, which was $V_0/2$, and in the direction to subtract from the original network voltage. At this instant, the voltage on the network drops to zero. The voltage at the open end of the network is shown as V_2 . At the same instant, all of the network elemental capacitances have been discharged to 1/2 of their original voltage, which means that they now store only 1/4 of their original energy. But the energy delivered to the load so far is only 1/2 of what it will be before the discharge cycle is complete. This is because it will continue until the network reflected wave has traveled back to the input end of the line for a total time duration of $2T$, twice the delay-time of the transmission line network. By that time, all of the elemental capacitances will have been completely discharged, the voltages on the network and across the load will drop to zero, and current flow in the discharge loop will terminate. The load energy, which is the product of load power and pulse duration, can be expressed as

$$\left(\frac{V_0}{2}\right)^2 \times 2T = \frac{V_0^2 T}{2Z_0}.$$

Load energy, W_L , at this time is equal to the amount originally stored in the network. At the halfway point, 3/4 of the original network energy had been removed from the elemental capacitances of the line, but only half of it had been delivered to the load. What happened to the other 1/4? At the halfway point—and *only* at the halfway point—the discharge current, which remains constant throughout the entire discharge cycle, is flowing through all of the elemental inductance of the line. The energy stored in the magnetic field of this inductance is equal to the missing amount, as shown in the following derivation:

$$Z_0 = \sqrt{L/C}$$

$$L = CZ_0^2$$

$$i = \frac{V_0}{2Z_0}$$

$$W_L = \frac{1}{2} Li^2 = \frac{1}{2} CZ_0^2 \times \left(\frac{V_0}{2Z_0}\right)^2 = \frac{CZ_0^2 V_0^2}{8Z_0^2} = \frac{1}{8} CV_0^2 = \frac{1}{4} W_0.$$

All of the network energy has been transferred to the load in the form of a pulse that is twice as long as the time-delay of the transmission-line PFN, which is what one might expect from a load matched to the characteristic impedance of a lossless source. Suppose, however, that the load is not matched to the source impedance.

Figure 9-5 shows the same transmission-line PFN connected to a load having resistance equal to half of the line characteristic impedance, or $R_L = Z_0/2$. This results in what is called a 2:1 overmatch condition. The voltage reflection coefficient, G , of this load is $-1/3$, as shown by




Figure 9-5. Discharge of transmission-line pulse-forming network into 2:1 overmatch.

$$\frac{\frac{Z_0}{2} - Z_0}{\frac{Z_0}{2} + Z_0} = -\frac{1}{3}$$

From conventional transmission-line theory, a designer would expect a mismatch or reflection loss equal to the square of the voltage reflection coefficient, or $1/9$. To see how (or if) this is true, the same circuit analysis will be followed as before.

When the discharge switch is closed at time t_0 , current will flow in the discharge circuit. It will be limited by the load resistance in series with the characteristic impedance of the network, whose total impedance is $1.5 Z_0$. The load voltage, therefore, will be negative-going as before and equal to the product of the discharge current and the load resistance, or $V_0/3$. The voltage across the input end of the network, V_1 , must be equal in voltage but opposite in polarity in order to satisfy the requirement that there be no voltage across the switch and diode combination. This is satisfied by the instantaneous discharge of the first increment of transmission-line capacitance by an amount that equals $2/3$ of the original network voltage. This incident wave propagates as before toward the open-circuited right-hand end of the transmission line, where it is once again reflected in the same direction and amplitude as it arrived, producing an instantaneous change in the voltage at the output end, V_2 , of $4/3 V_0$. This voltage charges up the endmost capacitive element in the opposite direction by the amount $V_0/3$. At this instant, all of the capacitive elements of the line have been discharged to $1/3$ of their original voltage, storing $1/9$ of the original energy. But only $4/9$ of the original energy has been delivered to the load so far, even though $8/9$ of the energy has been removed from the line capacitance. (Once again, we find that the unaccounted-for energy, equal to $4/9$ of the original energy, is stored in the line inductance.) When the $2/3 V_0$ reflected wave propagates back to the input of the line, after a delay of $2T$ from the beginning of the pulse, the voltage at the input, V_1 , also drops to $1/3 V_0$ in the reverse polarity. The discharge would continue were it not for the unidirectional property imposed by the diode in series with the switch.

In order for anything more to happen, the current would have to reverse. Because it is unable to, it stops. For the end condition, the network capacitance is charged to $1/3 V_0$ of the opposite polarity to the original charged condition. Evaluating the product of load power and pulse duration shows that an amount of energy has been delivered to the load equal to $8/9$ of the amount originally stored, and the remainder, the "mismatch loss," is not lost at all but is stored in the network capacitance. It is equal to $1/9$ of the original energy, as shown in the following derivations:

$$W_L = \frac{\left(\frac{1}{3}V_0\right)^2}{\frac{Z_0}{2}} \times 2T = \frac{4V_0^2}{9} \times \frac{T}{Z_0} = \frac{4}{9} CV_0^2 = \frac{8}{9} W_0$$




Figure 9-6. Discharge of transmission-line pulse-forming network into 2:1 undermatch.

$$W_{\text{stored}} = \frac{1}{2} C \frac{(V_0)^2}{3} = \frac{1}{9} \times \frac{1}{2} C V_0^2 = \frac{1}{9} W_0.$$

It could not have been "lost," because the only element in the circuit capable of power dissipation is the load resistance.

Figure 9-6 shows what happens if the sense of the mismatch is inverted, producing an overmatch condition of the same magnitude. The load resistance is now $R_L = 2Z_0$, and the total discharge-loop impedance is $3Z_0$. The load voltage is $2/3 V_0$ in the negative direction, so the network input voltage is also $2/3 V_0$ but in the positive direction. The $1/3 V_0$ incident wave travels down and back as before, but when it again reaches the input end of the network $2T$ after the start of the pulse, the voltage on the network is still positive, and so the current does not have to reverse in order for the discharge process to continue. The $1/3 V_0$

reflection from the open far end of the line is confronted with a re-reflection coefficient of $+1/3$ when it reaches the input end again. The coefficient is derived from

$$\frac{2Z_0 - Z_0}{2Z_0 + Z_0} = +\frac{1}{3}$$

And so a new incident wave is generated with a magnitude of $1/9 V_0$ that makes the round trip during the interval between $2T$ and $4T$, only to be supplanted by yet another wave with a magnitude of $1/27 V_0$, and so on. Each successive wave is one-third the size of its predecessor. The discharge, therefore, is a piece wise approximation of the overdamped double-exponential discharge of the simple L-C resonant circuit. (This is a variation of Zeno's paradox: how many jumps does it take to get to the end of a line if each jump is $1/3$ of the remaining distance?)

The undermatched case is one to be avoided. The half-control switch hangs




Figure 9-7. Discharge of 2:1 overmatched network with bidirectional switch.




Figure 9-8. Discharge of transmission-line network into short-circuit load.

up because it is not forced into a non-conducting state by the action of the external circuit, an effect referred to as load or circuit commutation of the switch in inverter and converter parlance. Compare this state to the matched and overmatched discharge cases discussed before, where the circuit forced the switch current to zero, or to less than a "holding" value of current, while simultaneously reducing the forward voltage either to zero or even slightly negative.

Noting that the undermatched case resembles a stair-step version of the overdamped R - L - C discharge circuit, it might be of interest to observe how the discharge could evolve in the overmatched case if the switch were made bi-directional (by removing the schematic diode in series with our perfect switch). This is illustrated in Fig. 9-7. The reflection coefficient seen by the network, looking toward the input end is $-1/3$, instead of $+1/3$ as in the undermatched

case, and the incident voltage wave is $2/3 V_0$, instead of $1/3 V_0$. Not surprisingly, the waveforms are now stair-step approximations of those of the oscillatory discharge of the underdamped $R-L-C$ circuit.

The last discharge case to be considered is of very great interest. It is with a short-circuited, or zero-impedance, load and represents a microwave tube that has experienced an internal high-voltage arc. The waveforms for this case are shown in Fig. 9-8. Even though the load impedance is zero, the discharge current is still limited by the characteristic impedance of the network. There is zero load voltage and zero input-end network voltage at the instant of switch closure.

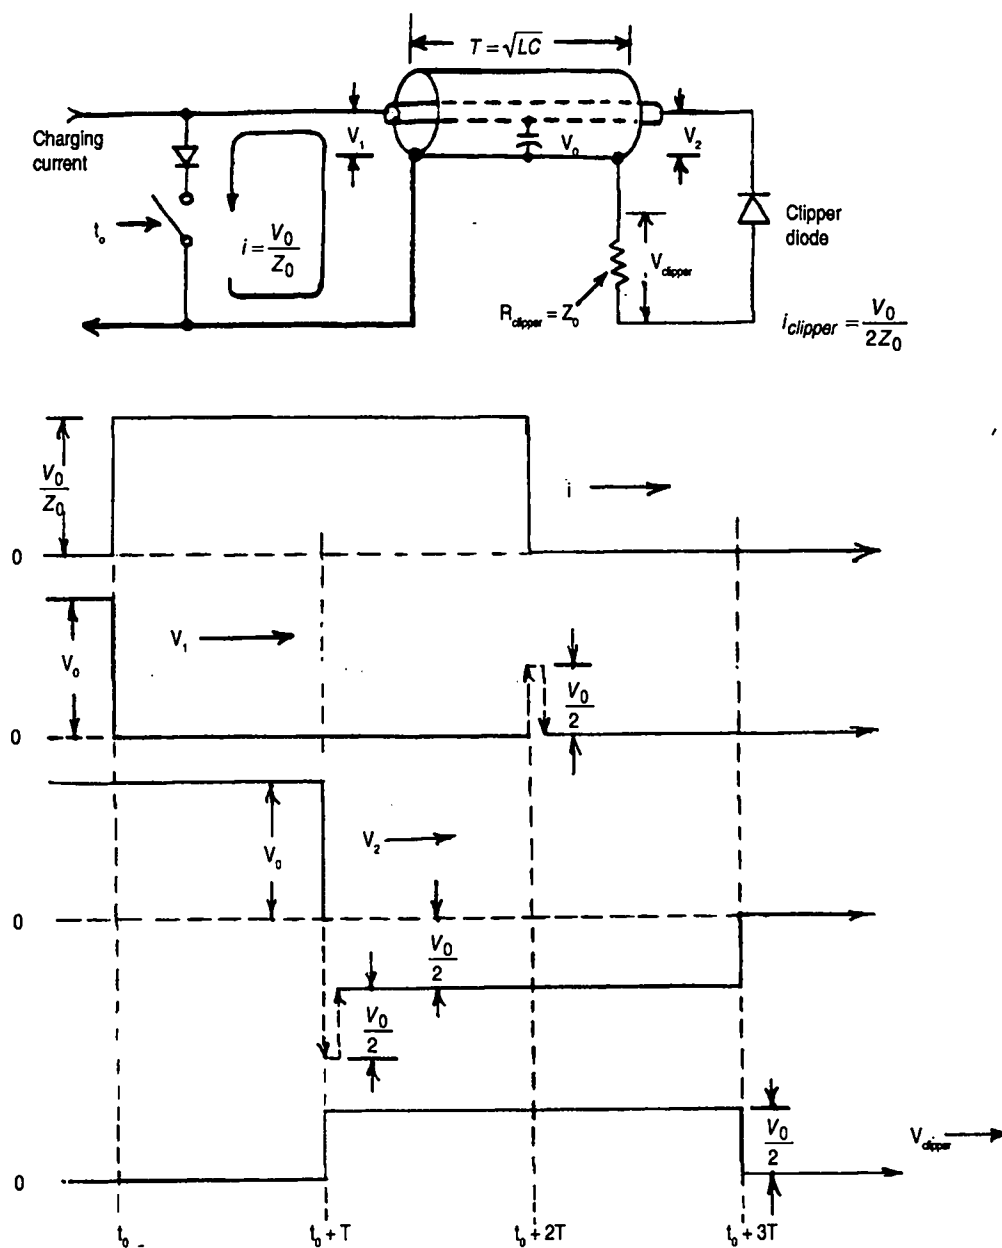


Figure 9-9. Discharge of transmission-line network with end-of-line clipper circuit into short-circuited load.

These conditions correspond to a rightward-traveling incident wave equal to the full network voltage. When the incident wave reaches the open end of the line, all of the network capacitances have been discharged, but the energy that had been stored in them is now stored in the network inductances, in the form of $1/2 LI^2$. When the incident wave is reflected from the open end of the line, the energy stored in the inductance begins being transferred back to the network capacitances, but its polarity is reversed as before, befitting an underdamped oscillatory condition. (Remember this when we later discuss a special kind of discharge circuit called the Blumlein network.) After twice the network delay-time has elapsed from the closing of the discharge switch, all of the network capacitances have been sequentially recharged in the reverse direction to the original network charge voltage, and the unidirectional switch blocks further current flow because it would have to reverse at this point. (This is, after all, the limiting case of overmatch.) No energy has been dissipated in our lossless ideal circuit, but what happens now? What can happen now will be discussed later when network charging circuits are considered. In general, however, the foregoing situation should be avoided for reasons having to do with objectionable switch and capacitor stresses. (One of the conditions capacitors—and not just polarized capacitors—dislike most is voltage reversal, and the faster it is done the less they like it.)

The back-swing clipper circuit illustrated in Fig. 9-9 prevents the previous shorted-load scenario from playing out. Shown as an end-of-line clipper, the circuit is connected directly across the open end of the PFN. It comprises a clipper diode and a load resistor that have a total forward-voltage-to-current ratio equal to the network characteristic impedance. For matched-load or undermatched conditions, there is no tendency for open-end network voltage to ever reverse polarity. As soon as load conditions tend toward the overmatched case, however, the open-end voltage will start to reverse. At that instant, the clipper diode circuit, formerly reverse-biased, is now forward-biased and conducts, introducing the clipper resistance into the circuit. For the shorted-load case illustrated, the conditions following switch closure, at least for the first delay interval, T , are identical with those of Fig. 9-8. At the precise instant when the voltage at the former open-circuited end tries to fully reverse, when the V_0 incident wave arrives, the negative-going reflected wave, instead of being fully reflected, encounters a situation just like the matched-load discharge case when the discharge switch is first closed, except the wave is upside-down.

As energy is transferred from the network inductance to the endmost capacitance increment, the current, which had been V_0/Z_0 , drops to $V_0/2Z_0$ because of the sudden appearance of the clipper-circuit resistance. The voltage across the clipper resistance is $V_0/2$ and positive-going in polarity. The network voltage simultaneously drops to $V_0/2$, also with a polarity opposite to that in the original matched-load case, thus satisfying the condition for no voltage drop across the perfect forward-biased clipper diode. (The dashed-line transitions shown in the figure occur simultaneously, so the network voltage never reaches the full reverse voltage indicated.) As in the matched-load case of Fig. 9-4, an incident wave of $V_0/2$ is launched, but this time from the output end back to the input end. When it gets there, after an interval of $2T$ from initial switch closure, it is confronted




Figure 9-10. Capacitor charging with resistive current limiting.

with a short-circuit, comprising the conducting switch and the short-circuited load in series that has a reflection coefficient of -1. This causes the inversion of the incident wave, which is consistent with having zero voltage across a short-circuit. When this reflected wave finally makes its way back to the clipper circuit, after a total interval of $3T$, all of the charge and energy originally stored in the network will have been delivered to the clipper resistance, and this particular pulse is over. In large, high-power pulsed transmitters, it is customary to monitor current in the clipper circuit, which is normally zero, and shut down subsequent pulse trigger signals if current is detected. If and when trigger signals should be resumed then becomes a decision left to the transmitter control logic or the operator.

9.1.2 The charging circuit

Having reviewed many aspects of network discharge, we will now consider aspects of the initial charging of the network and charging between successive pulses. Figure 9-10 shows the most basic way of charging a capacitance: from a source of both voltage and current, shown as a battery, V_{dc} . When the switch is closed at t_0 , current will flow from the battery through R and will eventually charge up C to the full value of V_{dc} , accumulating energy throughout the process. This charging current can be defined as

$$i(t) = \frac{V_{dc}}{R} e^{-t/RC}.$$

The history of voltage and current, starting at t_0 , is shown. Power flow is the product of voltage and current, and the accumulated energy, W_C , is the time integral of this product from zero to infinity,

$$\begin{aligned}
 W_c &= \int_0^{\infty} i(t) \times V_c(t) dt \\
 &= \int_0^{\infty} \frac{V_{dc}}{R} e^{-t/RC} \times V_{dc} (1 - e^{-t/RC}) dt = \frac{V_{dc}^2}{R} \int_0^{\infty} (e^{-2t/RC}) dt \\
 &= \frac{V_{dc}^2}{R} \left[-\frac{RC}{2} e^{-2t/RC} \right]_0^{\infty} = \frac{V_{dc}^2}{R} \left(RC - \frac{RC}{2} \right) = \frac{V_{dc}^2}{R} \times \frac{RC}{2} \\
 &= \frac{1}{2} CV_{dc}^2,
 \end{aligned}$$

which, as we see, yields the familiar result that the eventual stored energy is $1/2 CV^2$. Also shown is why charging this way is not the best of ideas. The energy removed from the source is the product of the charge transferred and the source voltage, where transferred charge is the time integral of the current from zero to infinity, or

$$\begin{aligned}
 V_{dc} \int_0^{\infty} i_{\text{charging}}(t) dt &= V_{dc} \int_0^{\infty} \frac{V_{dc}}{R} e^{-t/RC} dt = \frac{V_{dc}^2}{R} \left[-RC e^{-t/RC} \right]_0^{\infty} \\
 &= \frac{V_{dc}^2}{R} \times RC \\
 &= CV_{dc}^2,
 \end{aligned}$$

which is twice the energy stored in the capacitor. An amount equal to the stored energy was lost in the charging process, or

$$W_{\text{lost}} = CV_{dc}^2 - \frac{1}{2} CV_{dc}^2 = \frac{1}{2} CV_{dc}^2.$$

The energy was dissipated in the charging resistance. The ohmic value of that resistance does not affect this phenomenon. Remember, we will be striving to keep the source resistance as close to zero as possible, so alternative methods of limiting the capacitor-charging current must be found.

One such method of limiting network charging current with circuit elements other than resistance is shown in Fig. 9-11. Instead of resistance, a charging inductor, L_{ch} is used, and the resulting form of charging the capacitance of a pulse-forming network is called dc-resonant charging. The charging inductor and the network capacitance form a series-resonant circuit. (The primary winding of the output pulse transformer can be considered to present negligible impedance to the charging waveform.) At the end of a discharge event, when the network voltage has been reduced to zero (at time t_0 in the figure), the dc supply voltage, V_{dc} , is momentarily applied across the charging inductor. A sinusoidal current begins to flow because there is virtually no damping in the circuit. The average charge leaving the dc source is equal to the product of the charging current and the duration of the charging interval, or




Figure 9-11. DC-resonant charging of pulse-forming network.

$$Charge_{avg} = \frac{2}{\pi} \times \frac{V_{dc}}{\sqrt{\frac{L_{ch}}{C_n}}} \times \pi \sqrt{\frac{L_{ch}}{C_n}} = 2V_{dc} C_n.$$

Voltage across the network capacitance begins to build up with the wave shape of an inverted cosine wave. The period of the oscillatory current is

$$2\pi\sqrt{L_{ch}C_n},$$

where C_n is the network capacitance. After a charging interval equal in time to a half-period, charging current has reached zero and the network capacitance has been charged to twice the dc supply voltage. The diode in series with the charg-

ing inductor is called a charging diode. It prevents charging current from reversing polarity, which would discharge the network capacitance again. The charging current is limited in amplitude by the characteristic impedance of the series-resonant circuit,

$$\sqrt{\frac{L_{ch}}{L_n}}$$

As shown in Fig. 9-11, the energy stored in the network capacitance at the end of the charging interval is equal to the energy removed from the dc source during the interval:




Figure 9-12. DC-resonant charging with shorted load and no clipper circuit.

$$W_{\text{stored}} = \frac{1}{2} C_n \times (2V_{dc})^2 = 2C_n V_{dc}^2$$

$$W_{\text{removed}} = 2V_{dc} C_n V_{dc} = 2C_n V_{dc}^2$$

With this circuit arrangement our goals have been achieved; there is no resistive current-limiting and no energy loss. (It should be mentioned that an impression exists in some quarters that the purpose of dc-resonant charging is to double the dc power-supply voltage as it is applied to the network capacitance. This may or may not be a desirable attribute in itself. It is, however, a necessary by-product of obtaining total energy transfer from dc power source to network capacitance.)

DC-resonant charging can be the source of some interesting, if anomalous, behavior. This is especially the case for a simple discharge circuit having no clipper circuit when it is confronted with a short-circuited load, as shown in Fig. 9-12. The first of the discharge and recharge intervals illustrated are normal ones, with load impedance R_L equal to network characteristic impedance Z_0 . For some reason, however, at the beginning of the second discharge event the load experiences a short-circuit, such as an arc in a microwave-tube load. As discussed earlier, in the absence of a clipper circuit the network voltage reverses at the end of the two-way network delay, $2T$. At that instant, the charging inductor has a voltage across it of $3V_{dc}$ instead of the normal value, V_{dc} . Therefore, the cosinusoidal resonant ring-up of network voltage will carry the network voltage from its starting value of $-2V_{dc}$ through an excursion totaling $6V_{dc}$ to a final value of $+4V_{dc}$. It goes almost without saying that if the load did not arc on the previous pulse, it most assuredly will on this pulse. If it does, the result will be an end-of-pulse network voltage of $-4V_{dc}$ and an initial voltage across the charging inductor of $5V_{dc}$. The ring-up this time will be from the initial value of $-4V_{dc}$ to a final value of $+6V_{dc}$, an excursion of $10V_{dc}$. This process will continue, pulse after pulse, with the network charging each time to a value which is $2V_{dc}$ greater than the previous pulse until something in the transmitter overloads or breaks down. It is just such behavior that the clipper circuit prevents.

As we will see later, there are other switch-mode methods of charging PFN capacitance that share the "lossless" behavior of dc-resonant charging. (Lossless, in the sense that there will be no energy loss if circuit resistances can be reduced to zero. Any circuit that limits current with reactance rather than resistance falls in the "lossless" category.)

9.1.3 An illustrative example of a line-type modulator

Figure 9-13 shows a simplified illustration of a high-power line-type modulator application that represents an actual modulator used at Litton Industries' high-power test facility, which tests klystrons that can accept beam input power up to 90 MW peak. In this example, the load connected to the modulator by means of a step-up pulse transformer is a Varian VA-812 klystron (many of which are still in service in radar applications). Typical operating conditions are shown in a box in Fig. 9-13. At 250-kV beam voltage and 250-A peak-pulse beam current, the dynamic load impedance of the klystron is 1000 ohms. The dynamic resistance is not constant with beam voltage. It varies inversely with the square


root of it. Such high beam voltage is not practical to obtain directly from a PFN, especially because it would have to be charged to 500 kV in order to deliver 250 kV to a matched load.

A pulse transformer step-up ratio of 10:1 yields a practical transformer design and a manageable set of PFN and discharge-switch conditions. The impedance-transformation ratio of the pulse transformer is 10^2 , so that the load impedance referred to the primary side is $1000 \text{ ohms}/100 = 10 \text{ ohms}$, which is a very practical value for PFN characteristic impedance. Output-pulse transformer conditions are shown in another box, as are the PFN conditions. Note that in our example, which has no intentional resistance other than that represented by the klystron load, the energy stored in the network is 938 J, which is also the energy delivered to the klystron each pulse.

Heretofore, no mention has been made of the operating conditions of the discharge switch, shown here as a hydrogen thyatron. Another box in Fig. 9-13 shows these conditions. The network is charged to 50 kV, which is twice the 25-kV matched-load pulse at the pulse-transformer primary. The peak switch current, therefore, is 50 kV divided by twice the network characteristic impedance, or $50 \text{ kV}/(2 \times 10 \text{ ohms})$, or 2500 A. The average thyatron current is the peak current multiplied by the duty factor, 0.005, or 12.5 A. The average current is the component that is important in determining the dissipation in the ionized hydrogen, which is the conducting medium of the switch. Average current is used because the voltage drop across the ionized hydrogen is non-linear with current. In fact, voltage is nearly constant for current values around the nominal operating point. The RMS current is also important because it determines the dissipation in the ohmic components of the thyatron, such as the metallic current-carrying parts. (Note that the RMS current is considerably larger than the average current.) There is a third component of stress that is proportional the square of the operating voltage, to the magnitude of stray shunt capacitance associated with the switch, and to the pulse repetition frequency (PRF), irrespective of the duty factor. This is the "switching loss", and thyatrons are individually rated in terms of how much of it they can tolerate. There are external circuit means to limit the capacitive discharge current and support the switching voltage until the hydrogen is fully ionized. These are referred to as "magnetic assists." They use saturating magnetic components that can greatly reduce this stress.

Another box of Fig. 9-13 lists the charging-circuit conditions. The charging inductor has been sized so that the charging current just reaches zero before the next discharge trigger is applied to the discharge switch at the highest PRF of 333 pps. The average charging current, which determines the average power removed from the dc power supply, is 12.6 A. This current, when multiplied by the 25-kV supply voltage, gives the average dc power input of 313 kW. Note that this is identical with the average klystron input power (62.5 MW peak power \times 0.005 duty factor, or 313 kW average power). The RMS value of the charging current is also of interest because it determines what ohmic losses there will be in the charging-inductor wire. If this wire is big enough in cross-sectional area, the loss component can be made negligibly small. In fact, it has been assumed throughout this illustrative example that ohmic, magnetic-core, ionized-plasma, solid-state-junction, and dielectric losses are not only negligible but are zero, so

Charging conditions	Pulse-forming network conditions
<ul style="list-style-type: none"> • Pulse repetition interval (PRI) = 3 ms = $\pi\sqrt{L_{Ch}C_N}$ • $L_{Ch} = \frac{PRI^2}{\pi^2} \times C_N = \frac{3 \times 10^{-3}}{\pi^2} \times 0.75 \times 10^{-6} = 1.2\text{H}$ • $V_{dc} = \frac{V_N}{2} = \frac{50\text{kV}}{2} = 25\text{kVdc}$ • Peak charging current = $\frac{V_{dc}}{\sqrt{\frac{L_{Ch}}{C_N}}} = \frac{25 \times 10^3}{\sqrt{\frac{1.2\text{H}}{0.75 \times 10^{-6}\text{F}}}} = 20\text{A}$ • Average charging current = $20\text{A} \times \frac{2}{\pi} = 12.6\text{A}$ • DC input power = 25 kVdc x 12.6 A = 313 kW • RMS charging current = $1.11 \times I_{AVG} = 1.11 \times 12.6 = 14\text{A}$ 	<ul style="list-style-type: none"> • Network delay = Pulse duration/2 = $15\mu\text{s}/2 = 7.5\mu\text{s} = \sqrt{L_N C_N}$ • $Z_0 = 10\text{ ohms} = \sqrt{\frac{L_N}{C_N}}$ • $\frac{L_N}{C_N} = Z_0^2 = 10^2 = 100$ $L_N = 100 C_N$ • $\sqrt{L_N C_N} = \sqrt{100 C_N \times C_N} = 7.5\mu\text{s}$ $10 C_N = 7.5\mu\text{s}; C_N = 0.75\mu\text{F}$ • $L_N = 100 C_N = 75\mu\text{H}$ • $V_N = 2 \times 25\text{ kV} = 50\text{ kV}$ • Network energy = $\frac{1}{2} C_N V_N^2 = \frac{1}{2} 0.75 \times 10^{-6} (50 \times 10^3)^2 = 938\text{J}$



Discharge switch conditions	Output-pulse-transformer conditions	Klystron beam pulse conditions
<ul style="list-style-type: none"> • Peak switch current, I_{pk}, $\frac{V_N}{2Z_0} = \frac{50\text{kV}}{2 \times 10\text{ohms}} = 2500\text{A}$ • Average switch current = $I_{pk} \times \text{Duty factor} = 2500\text{A} \times 0.005 = 12.5\text{A}$ • RMS switch current $= I_{pk} \sqrt{DF} = 2500\text{A} \sqrt{0.005} = 177\text{A}$ 	<ul style="list-style-type: none"> • Secondary pulse voltage = 250 kV • Primary pulse voltage = 25 kV • Transformer turns ratio = N $\frac{V_{sec}}{V_{pri}} = \frac{250\text{kV}}{25\text{kV}} = 10:1$ • Impedance transformation = $N^2 = 100:1$ • Secondary load impedance = 1000 ohms • Reflected primary load impedance = $\frac{R_{sec}}{100} = \frac{1000}{100} = 10\text{ ohms}$ 	<ul style="list-style-type: none"> • Peak beam voltage = 250 kV • Peak beam current = 250 A • Peak beam power = 250 kV x 250 A = 62.5 MW • Pulse duration, T, = 15μs • Pulse repetition rate, PRR, = 333 pps • Repetition interval = 1/PRR = 1/333 = 0.003 s • Duty factor, DF, = Pulse duration x PRR = 15μs x 333 pps = 0.005 • Average beam power = Peak power x DF = 62.5 MW x 0.005 = 313 kW • Beam energy/pulse = Peak power x pulse duration = 62.5 MW x 15μs = 938 J

Figure 9-13. An illustrative example of a high-power, line-type modulator.

that the efficiency, from ac input to klystron beam pulse out, is 100%. (This is what one should expect from a true switch-mode circuit, which this one certainly is because its only active element, the thyatron, is clearly operated as a switch.)

How close can a physically realizable version of this circuit come to 100% efficiency? To give us an idea, the most critical component, the hydrogen thyatron switch, can have a conduction-interval efficiency of 99%. (For a perfect conduction-interval efficiency, the ratio of hold-off voltage to conduction-voltage drop would equal 100.) The losses in everything else, except for the solid-state charging and ac-dc rectifier diodes, depend primarily on how physically large they can be. (Minimum size and maximum efficiency are usually mutually exclusive attributes.) Overall efficiency exceeding 90% is quite achievable.




Figure 9-14. Discharge of the matched-load Blumlein circuit.

9.1.4 The Blumlein discharge circuit

In defending the use of an output-pulse transformer in the previous example, it was mentioned that getting high-amplitude output pulses directly from a network in the matched-load case requires that the network be charged initially to twice the output voltage, which has certainly been the case in all of the foregoing discussions. There are times, however, when a direct network-output-pulse amplitude must be as large as possible. These situations are usually due to limits in switch voltage because the switch must withstand the network charge voltage. The Blumlein circuit, one of a class of networks that is charged in parallel but discharged in series, uses two pulse-forming networks, as shown in Fig. 9-14. When the single discharge switch at the far left of the schematic diagram is open, the two networks, which have identical characteristic impedances and delay-times, are charged in parallel to the same voltage, V_N . Note that the network to the left is connected to the pulse-current return bus (ground) and, therefore, bears a strong resemblance to a conventional PFN operating into a short-circuited load. The PFN to the right is connected to the load in more conventional fashion, except the load resistance for the "matched-load" condition is $2Z_0$, instead of Z_0 .

When the discharge switch is closed at time t_0 , the instantaneous voltage at the left-hand end of the first PFN changes from V_N to zero, which is the same as launching an incident wave of amplitude V_N in the rightward direction (as was previously discussed in the case of the conventional PFN circuit with a short-circuited load). When this incident wave reaches the output end of the network after time T , it will attempt to reverse the voltage on the network. In attempting to do this, however, it is forced to take notice of the fact that its output end is not open-circuited as before. Instead, it is connected to a second network of impedance Z_0 in series with a load resistance of $2Z_0$.

For the split nanosecond before current can flow, a situation exists where a voltage of $2V_N$ ($+V_N$ across the right-hand network and $-V_N$ momentarily across the left-hand network) is in series with a circuit comprising twice the individual network impedance Z_0 plus the load impedance $2Z_0$ for a total impedance of $4Z_0$. The load-loop peak current is $2V_N/4Z_0$, or $V_N/2Z_0$, and the load voltage is the load current multiplied by the load impedance, $2Z_0$, or $V_N/2Z_0 \times 2Z_0 = V_N$. The load voltage is equal to the network charging voltage. The voltage across each network instantaneously drops to 1/2 of what it was at the instant of input-network voltage reversal, with the equivalent of $V_N/2$ incident waves propagating to the left in the left network and to the right in the right network. The left-going wave, after another interval of T , reaches the input end and the conducting switch, which has a voltage reflection coefficient of -1. This inverts the wave (maintaining the criterion of no voltage across a short-circuit) and sends it back to the right. Once the wave departs in this direction, however, there is no more switch current. The switch, therefore, conducts a current V_N/Z_0 for a time interval of $2T$ for a total charge transport of $2T$ times V_N/Z_0 . The charge stored in each network is $V_N C$, but because $C = T/Z_0$, the total charge stored is $2T \times V_N/Z_0$, which is the same as that handled by the switch. The energy stored in each network is the familiar $1/2 CV_N^2$, for a total of CV_N^2 . The load energy is

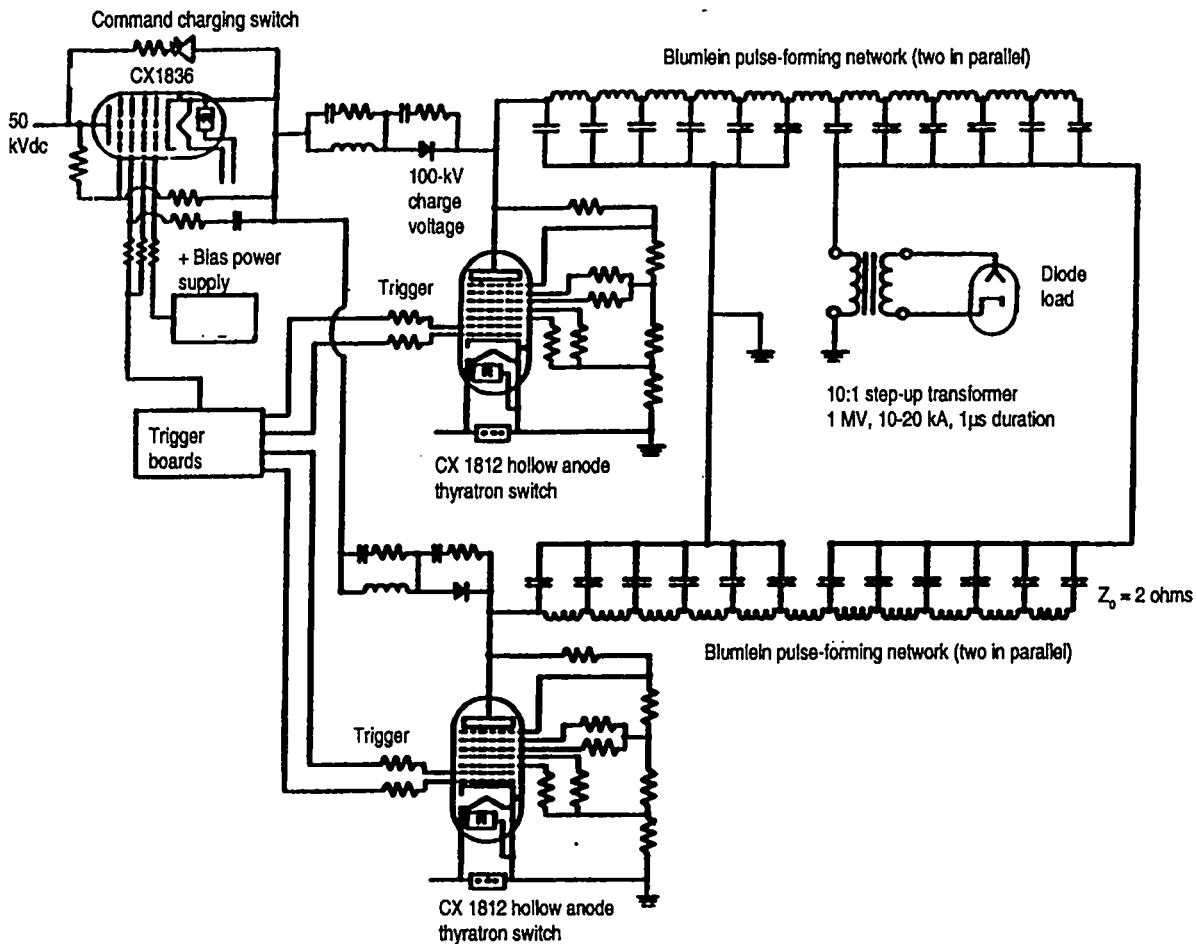


Figure 9-15. The BANSHEE, a megavolt-output Blumlein-type pulse modulator.

$$\frac{V_N^2}{2Z_0} \times 2T,$$

where $T = CZ_0$, or CV_N^2 , which is the same as the original stored energy.

Why would anyone go to all of this trouble just to double the output from the network when a pulse transformer can do the same thing, and more? Let's see. Figure 9-15 is a simplified schematic diagram of the Beam Accelerator for a New Source of High-Energy Electrons (BANSHEE) line-type modulator. It was designed to produce an output-pulse amplitude of 1 MV at a peak-pulse current of 10,000 A into one of a number of exotic super-high-power microwave devices, including a relativistic klystron, a virtual-cathode oscillator (or Vircator), and a large-orbit gyrotron. The practical upper limit of a pulse-transformer step-up ratio is on the order of 10:1, which means that the primary voltage needed for a 1-MV peak would have to be 100-kV peak. To get this kind of voltage from a conventional PFN pulser, the network-charging voltage and switch hold-off voltage would have to be 200 kV, which is a lot. In this case, it is too much for practical discharge switches. Note in the figure that the complete pulser is a parallel/parallel arrangement with two complete pulsers. Each pulser has its

own discharge thyratron-switch tube in parallel, and each of the pulsers has two network pairs, also in parallel, for a total of eight networks for the complete modulator.

During the discharge process, as shown in Fig. 9-14, the voltage on the left-hand (grounded) network reverses by an amount equal to half of the original charge voltage. This is an extremely stressful experience for capacitors, meaning that the capacitance of the first network is more likely to experience failure than that of the second. In the Blumlein circuit, two networks are charged and discharged. (There are circuit arrangements, as might be expected, that permit larger numbers of networks to be charged in parallel and discharged in series, but they result in even greater voltage reversals.)

Another practical manifestation of the Blumlein connection is the water-filled Blumlein, illustrated in Fig. 9-16. High-purity water is a very poor conductor but has a very high relative permittivity of 80, so that a transmission line using water as the dielectric will have low characteristic impedance and relatively high dc capacitance. The water-filled Blumlein illustrated uses three coaxial conductors. The intermediate conductor serves as the outer conductor for an inner coaxial line and the inner conductor for an outer one. Constructed with the dimensions




Figure 9-16. The water-filled coaxial short-pulse, high-power Blumlein network.




Figure 9-17. Schematic representation of the Marx-type impulse generator.

shown, the characteristic impedances of both lines were approximately 5 ohms. When filled with high-purity water, the dc capacitance was $0.01 \mu\text{F}$ and the propagation delay was 23 ns. With a 10-ohm load (twice the individual characteristic impedances) connected between innermost and outermost conductors and a charge voltage of 250 kV applied to the intermediate conductor—rapidly, because the inherent discharge time-constant is $100 \mu\text{s}$ —the Blumlein is ready to fire. This is accomplished by means of a high-voltage trigger signal applied to an annular “coaxial midplane” electrode. It causes an ionized conduction path between the intermediate and outer conductors. This sets in motion the discharge cycle as described in Fig. 9-14, and after a 23-ns delay, a 46-ns output pulse across the 10-ohm load at 250 kV peak voltage is produced. The total energy stored in the coaxial networks is 310 J, which is not much, but when it is released over an interval of only 46 ns, the peak power is 6740 MW.

9.1.5 The Marx impulse generator

Perhaps the best known of the high-voltage pulse generators that charge energy stores in parallel and discharge them in series is the Marx impulse generator. The "networks" are seldom more than capacitors in series with current-shaping resistors. In Fig. 9-17 the capacitors are labeled C_1 through C_N and the current-shaping resistances are R_F . But practical generators have been built with a large number of stages to produce "erected" voltages of hundreds of kilovolts from a charging voltage of 10 kV or so. The Marx is erected by means of spark gaps after the individual capacitors have been charged through isolating resistors, R_{CH} . (Note that these resistors do not have to determine the charging current, so that charging efficiency need not be limited to 50%.) Only the bottom-most gap needs to be triggered in most situations because subsequent gaps are sequentially overvoltaged by the effect of the stray capacitances to ground, shown as C_S . Erection times are measured in nanoseconds. A Marx generator is often used to charge up a water-filled Blumlein pulser because it needs to be done in a hurry.


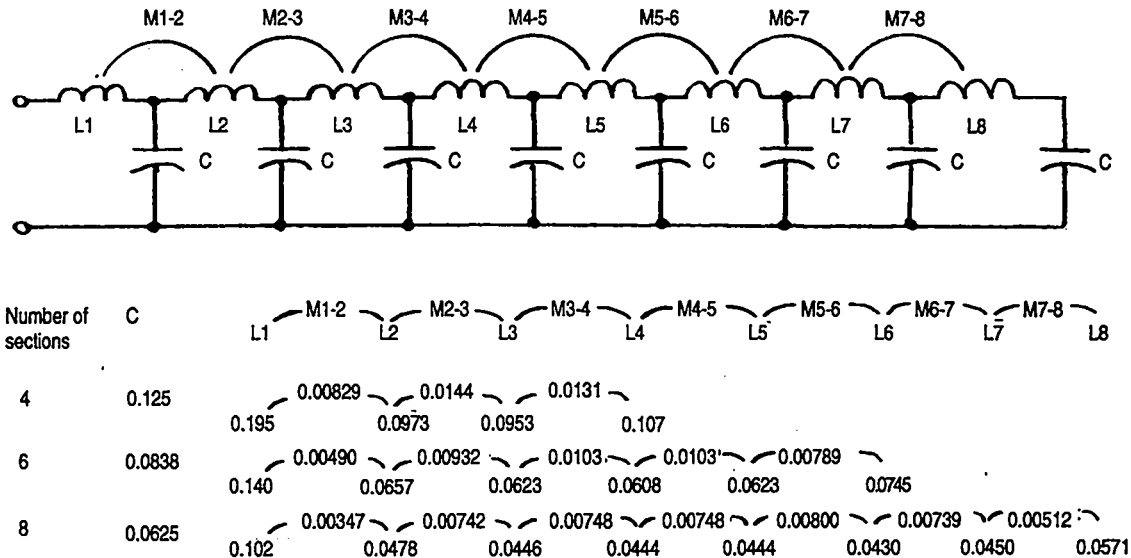


Figure 9-18. The Guillemin lumped-circuit pulse-forming network types.



Inductances in μH , capacitances in μF for $Z_0 = 1 \text{ ohm}$ and $\tau = 1 \mu\text{s}$.
For other values, multiply inductance values by $Z \times \tau$ and capacitance values by τ/Z .

Figure 9-19. The Guillemin Type E pulse-forming network.

9.1.6 Pulse-forming networks

As has been mentioned before, the ideal PFN is a distributed constant transmission line that has no finite filter cut-off frequency, a characteristic impedance equal to the transformed load impedance, and a propagation delay equal to half of the desired output-pulse duration. However, with propagation velocities of approximately 1 ft/ns, such networks are hardly practical, except when pulse durations are well below the submicrosecond range. Therefore, lumped constant artificial transmission-line segments are used to satisfy most requirements. As the inductance and capacitance per section (and the total number of sections) become finite, the waveforms produced by discharging such line segments have increasing amounts of pulse-top ripple, especially near the beginning and end of each pulse. This is due to the fact that as the number of network sections becomes finite instead of infinite, the incremental inductance and capacitance values become finite instead of infinitesimal. As the number of network sections gets smaller, the more pronounced the effect becomes. To mitigate this effect, Ernst Guillemin, who gave his name to the commonly used networks, postulated that networks be synthesized with intentionally finite rise-and-fall intervals and with non-rectangular wave shapes, such as trapezoidal and parabolic (because they are what you will end up with, anyway). Moreover, he postulated networks that are not just transmission-line equivalents. Figure 9-18 shows the topology of the Guillemin network types, normalized to five sections each.

The Type A network is a series-connected cascade of parallel-resonant circuits, with a series input capacitor and output inductor. The individual parallel-resonant sections have resonant frequencies that are the next highest odd harmonic of the preceding one. (The fundamental frequency corresponds to basic pulse duration, so we are talking about harmonics 3, 5, 7, etc.) The network,




Fig. 9-20. Computer-derived wave shapes for five-element Type C networks (fundamental and third, fifth, seventh and ninth harmonics).

therefore, emulates the Fourier series for a rectangular wave shape, which is the fundamental frequency and the higher-order odd harmonics.

The Type B network is a transmission-line equivalent, but its element values are skewed as a result of the deliberate non-rectangular-waveform synthesis.

The Type C network, comprising parallel-connected series-resonant circuits also of ascending odd-harmonic resonant frequency, is the dual of the Type A network.

The Type D network is another transmission-line synthesis. But it has the highly desirable attribute of possessing equal-value capacitors—desirable because capacitors are the more difficult of the two circuit elements to physically realize. However, the network accommodates them by using series negative inductance, which is even more difficult to realize.

Figure 9-19 shows the Type E network, which is by far the most commonly used. Note how the negative inductance of the Type D network has been realized by making use of mutual inductance between the individual network inductance segments. A properly wound solenoid with periodic tap points for capacitor connection can be made to have the necessary values of mutual and self-




Figure 9-21. A high-performance 10-section Type E pulse-forming network.

inductances. Returning to Fig. 9-18, the Type F network is seen to be the dual of the Type B network.

Figure 9-20 shows how computer analysis would predict the wave shapes from five-section, Type-C networks for four different synthesis criteria: trapezoidal waveforms for 5% and 10% rise-and-fall intervals; and parabolic waveforms for 5% and 10% rise-and-fall intervals. An example of a high-performance, 10-section network is shown in Fig. 9-21. It is capable of producing 1.3- μ s-duration pulses of 13 MW peak power with 0.25- μ s rise time. The coil is 22 in. long and 2.5 in. in diameter with a pitch of 1.7 turns/in. The capacitors are charged to 26 kV and have a capacitance of 0.047 μ F.

Not all practical pulse-forming networks have evolved from Guillemin's work, however. Many high-performance pulser networks, sometimes referred to as Rayleigh networks, use multiple sections of nominally identical inductive and capacitive elements. In these cases, the individual inductances are magnetically isolated either by shielding, relative physical orientation, or both. In addition, they are separately adjustable in value, often by means of motorized-slug trim-

mers, so that adjustments can be made with the pulser at full operating conditions. The adjustments are sometimes made to trim characteristic impedance and/or pulse duration, but more often they serve the purpose of minimizing the pulse-top ripple amplitude. Although controlling this ripple is important in almost all applications, it is critical in some.

These precision adjustments must be made at high level while the network components are subjected to full stress. Both capacitors and inductors are subjected to mechanical stresses that tend to change their dimensions: capacitors because of the voltage across them, and inductors because of the current through them. In fact, the magnetic force on an inductor tends to distort it physically in such a way as to increase its inductance. (All natural systems tend to gravitate, so to speak, toward the lowest energy state. Nature assumes that, in a current-carrying system, the product of LI will somehow remain constant, which is the dual of the constant CV product in an electrostatic system. If the force due to the flow of current in an inductor increases the inductance, it is assumed that there will be a corresponding decrease in circuit current. Stored energy, however, is $1/2 LI^2$, which will decrease.) Mechanical bracing to resist these forces is essential for both capacitor plates and, even more importantly, for the inductor winding. The force on a pulse-by-pulse basis will be manifested as an impulse proportional to the integral of $I^2 dt$, which is also the adiabatic heating action integral. Most PFN inductors are wound as single-layer solenoids. Magnetic force will tend to make such windings shorter and their diameters larger, the directions that yield even greater inductance. These are not simple forces to brace against. The physical shape for an inductor that is easiest to brace against magnetic-force deformation is the flat spiral, but it is by no means the easiest to integrate into a complete network topology.

Failure to adequately constrain physical deformation will lead to network performance that can vary excessively as the operating level is changed. This can pose a thorny problem when precision adjustments are required to achieve pulse-top flatness. These adjustments are often the nulling type, where a small value is obtained by subtracting two relatively large values. Unless both values change together and by the same amount, the small difference will no longer be small. (Note that performance obtained by computer analysis of network values can often be breathtaking, and neither magnetic nor electrostatic forces need alter them.)

Even if peak-pulse operating levels are invariant, a change in PRF can have the same, or worse, effect. This will be true if PRF should coincide with a mechanical resonant frequency of the physical structure, or a submultiple of the frequency.

In a line-type pulser, it is the nature of the line itself, or PFN, that determines all of the important attributes of the pulser output. This is both the great strength of this pulser and its great weakness.

9.1.7 Output-pulse transformers

Very few line-type pulsers are ever directly coupled to their loads—the most likely exception being the Blumlein type. In the great majority of cases, the load will be matched to the pulser by means of a step-up, pulse-optimized output




Figure 9-22. Pulse transformer characteristics and how they affect pulse wave shapes.

transformer. And the great majority of these transformers will have an unmistakable family resemblance to all other transformers. They will comprise a magnetic core over which the primary and secondary winding are usually wound. The secondary winding is wound as close to the primary winding as is consistent with high-voltage isolation of the high-end of the secondary winding. The farther away the secondary winding is, the greater the leakage inductance will be, which, along with the shunt capacitance of the windings and load, will determine the rise-time performance of the transformer. The other component of transformer inductance is the magnetizing inductance, which is in shunt with the load. It determines the droop, or voltage decrement, of the flat-top portion of the pulse.

Figure 9-22 shows a simplified equivalent circuit with all circuit values normalized to either the primary or secondary side of the transformer. The family of curves at the lower left shows the effect on output-pulse rise time of different damping factors (δ) where the load resistance is equal to source resistance ($R1 = R2$), which is usually the intended case for a line-type pulser. The damping factor relates load (or source) resistance to the combination of characteristic impedance of the leakage-inductance, $L1$, and shunt-capacitance, C , as

$$\delta = \frac{1}{2} \sqrt{\frac{R2}{R1 + R2}} \left(R1 \sqrt{\frac{C}{L1}} + \frac{1}{R2} \sqrt{\frac{L1}{C}} \right).$$

The other family of curves is for source a resistance of zero, where damping is provided only by the load resistance. This condition can be expressed as

$$\delta = \frac{1}{2R2} \sqrt{\frac{L1}{C}}.$$

As soon as there is voltage across the transformer windings, a component of current will begin to build up in the magnetizing inductance, $L2$. Assuming a rise time that is very short compared to pulse duration, the current can be expected to have a slope, di/dt , which is equal to $V/L2$. This current will develop voltage across the internal generator impedance, which will subtract from the generator voltage to produce linear droop in output voltage, as shown by the




Figure 9-23. A typical high-voltage, high-power output-pulse transformer.

right-most set of graphs in Fig. 9-22. (If the generator impedance is zero there will be no output voltage droop no matter how long the pulse or how low the magnetizing inductance. But both magnetizing and generator currents will tend toward infinity.)

The output-pulse transformer must be designed to handle the high voltages, up to 1 MV peak, as previously described. A typical high-performance transformer design is shown in Fig. 9-23. The secondary takes the form of parallel-connected, flat-sided conical windings, with the narrow ends at the low-side where the voltage differences between primary and secondary are near zero. As the voltage difference between primary and secondary increases with winding height, so does the spacing between windings. Leakage inductance is related to the spacing between windings and will be roughly half as great as it would be if the secondary had the same spacing all the way up. A transformer of this type is designed to be operated in insulating oil, for both its dielectric strength and heat-transfer property.

Very often the secondary will use a bifilar winding, the ultimate form of which is a coaxial cable. The bifilar winding is used to inject the filament excitation into the microwave tube whose cathode is connected to the secondary high side. Because the two closely coupled wires will have the same pulse voltage induced in them, whatever filament voltage is impressed on the wires at the low, or ground, side will also be between them at the high side. To minimize the required ampacity of the bifilar winding, a step-down filament transformer is often connected between high-end bifilar conductors and the tube heater leads, requiring, of course, that the bifilar winding leads be insulated for the higher transmission voltage. (It is usually easier and cheaper to design for the necessary insulation rather than design for larger-diameter wire.)

When pulse duration is long, core saturation is always a problem. The applied voltage across the transformer is unidirectional, so the magnetizing current is direct current. To maximize the volt-time performance of a core, dc bias is often employed. It is either injected into the primary winding through a bias choke, which isolates the pulse primary voltage from the bias-current dc power




Figure 9-24. What happens to energy stored in pulse-transformer magnetizing inductance.

supply, or into a tertiary winding on the transformer. The optimum amount of core-bias current will vary with transformer operating conditions. But the idea is always to swing the core magnetization over the region of its B/H curve, with the greatest slope corresponding to highest shunt inductance.

The current that builds up in the transformer magnetizing inductance produces an increasing amount of stored energy that grows as the square of that current. When the pulse ends, the energy stored cannot be transferred to the microwave-tube secondary load, because it is a diode and the magnetizing current is going in the wrong direction. It is, however, going in the right direction to continue through the modulator switch and the PFN. The network current due to magnetizing current is in the opposite direction to the normal network charging current, so it will tend to charge up the network in the reverse direction. Figure 9-24 shows what happens. The magnetizing current can be defined as

$$i_M = \frac{V_0 t}{L_M}$$

and magnetizing current at the end of a pulse, I_M , can be defined as

$$I_M = V_0 \times \frac{2T}{L_M} = \frac{V_0 2\sqrt{L_N C_N}}{L_M}$$

When current finally stops, the network will have a voltage across it in the reverse direction such that the energy stored in it will equal the energy that had been stored in the transformer magnetizing inductance at the end of the pulse. At the end of the pulse, the energy stored in magnetizing inductance, W_M , will be equal to

$$\frac{1}{2} L_M I_M^2 = \frac{1}{2} L_M \left(V_0 \times 2 \frac{\sqrt{L_N C_N}}{L_M} \right)^2 = \frac{1}{2} L_M \frac{(4V_0^2 L_N C_N)}{L_M^2} = 2 \frac{V_0^2 L_N C_N}{L_M}$$

After the pulse, the magnetizing current continues to flow through the network and switch, charging up the network in the opposite polarity from normal charging current until the energy stored in the magnetizing inductance is transferred to the network as reverse network voltage, V_{NR} , with energy W_{NR} .

$$W_{NR} = W_M$$

$$\frac{1}{2} C_N V_{NR}^2 = \frac{2V_0^2 L_N C_N}{L_M}$$

But because

$$V_0 = \frac{V_N}{2}$$

and

$$-\frac{1}{2}C_N V_{NR}^2 = \frac{1}{2} \frac{V_N^2 L_N C_N}{L_M},$$

then

$$V_{NR}^2 = \frac{V_N^2 L_N}{L_M}$$

$$\frac{V_{NR}^2}{V_N^2} = \frac{L_N}{L_M}.$$

Solving such an equation shows that the ratio of network reverse voltage to normal forward voltage, V_{NR}/V_N , is equal to the square root of the inductance ratio, L_N/L_M . This voltage-reversal effect is often used in line-type pulser design to provide the small voltage reversal that can encourage recovery of the modulator switch to its non-conducting original state.

9.2 Why do we care about pulse-top flatness?

There are as many answers to this question as there are engineers—or customers. Most engineers, for instance, really like to see flat-topped rectangular pulses on their oscilloscopes. But so far, we have been dealing with a type of pulse generator that inherently produces pulse-top ripple, which is something that can be problematic. This raises a question: at what point, and after how much circuit refinement and fine-tuning, must this entire category of pulse modulator be abandoned in favor of a type that gives better pulse-top performance?

To gain an understanding of this problem, let's first consider a type of modulation for a high-power RF amplifier that can produce the same information content in its RF output as that produced by a very short pulse, which is what line-type pulse modulators generate. Radar systems require short pulses (or their informational equivalents) in order to have range-resolving capability, which is the ability to clearly distinguish between individual targets, or "scatterers," that are closely spaced in range. The range-extent of transmitted pulses must be shorter than the target spacing so that individual, separate returns will come back from each target. We don't want the returns to be all smeared together. The RF bandwidth of the transmission system is related to $1/T$, where T is the pulse duration. Therefore, a 1-ns transmitted pulse, which is capable of approximately one-foot range resolution, requires a 1000 MHz bandwidth. Presumably, the equivalent information could be imparted by a system that used 1000-MHz instantaneous bandwidth, whether or not it put out a short pulse. This is what linear-chirp frequency modulation does. If, over an actual transmitted pulse

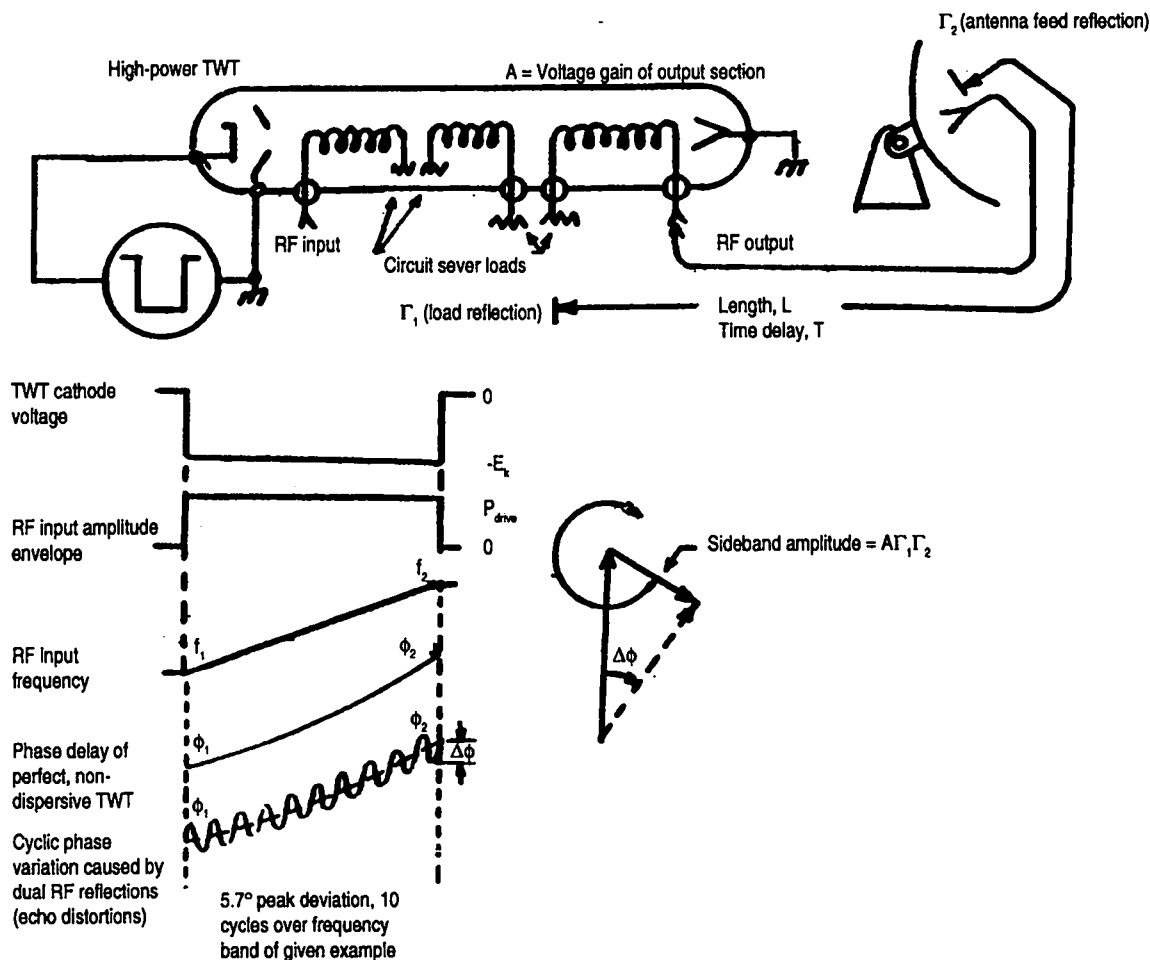


Figure 9-25. Cyclic phase deviation as a function of frequency due to phase-pulling of TWT.

duration of arbitrary length, the instantaneous frequency is linearly ramped by 1000 MHz, a frequency-decoding signal processor at the receiving end should be able to reconstruct a pulse of 1-ns-equivalent duration. What is transmitted is called an expanded pulse, and what comes out of the signal processor is called a compressed pulse. In an ideal system, the energy content of the long, received, expanded pulse and its short, compressed offspring is the same.

Before considering what effect pulse-top ripple would have on the output pulse in such a system, let us consider the general case of distortion in such a system. Figure 9-25 shows a transmitter using a broad-band TWT as high-power RF amplifier. It is connected to an imperfectly matched antenna feed through a length of transmission line. The TWT has three individual gain sections that are circuit-isolated by means of severs in the circuit, each resistively terminated. (Coupling between isolated circuit segments is by velocity and density RF modulation of the TWT beam.) The important part of the circuit of the TWT for this discussion is the output segment and its input termination, which is also imperfectly matched. (This is usually the case in actuality, too.) The voltage gain of the output segment in the forward direction is A , and the attenuation of an RF signal in the reverse direction is assumed to be negligibly small.

We have now established the conditions for transmitting echoes, which, on

reception, will be indistinguishable from true target returns. The generation of echoes in the short-pulse case can be fairly easily visualized. Using an example bandwidth of 500 MHz, the equivalent short pulse would be 2 ns in duration and would represent approximately 2 ft in range extent. A 3-m transmission line between TWT output and antenna feed input will have a delay each way of 10 ns. The antenna feed and the TWT circuit termination both have voltage-reflection coefficients. If we assume that they are both equal to 0.1, a transmitted pulse, upon reaching the antenna feed, will have a 10% voltage reflection, which will propagate back toward the transmitter, back into the TWT output, back down the output circuit segment, and finally reach the circuit termination, which will re-reflect 10% of the reflected signal, heading it back toward the antenna. Passing through the TWT output section, however, it will be amplified by a factor of A . If A is equal to 10, the echo, or re-reflected signal, which is called a "time-sidelobe" in signal-processing parlance, will have a voltage of $0.1 \times 0.1 \times 10$, or 0.1 times the original signal and will be delayed from it by twice the delay time between TWT output and feed input, or 2×10 ns (or 20 ns). And 20 ns after that, there will be another echo that will be 10% of the one before, or 1% of the original pulse. This phenomenon will repeat itself every 20 ns, with each successive echo being 10% of the one before, or smaller by 20 dB. (If this behavior seems reminiscent of time-domain reflectometry, it's because that's what it is.)

If we can replace the short transmitted pulses with a long transmitted pulse while maintaining the same instantaneous bandwidth and the same range resolution, it stands to reason that the same transmission-system distortion (echo distortion) will produce the same type of compressed-pulse, processed-signal distortion. Explaining why this is so intuitively rather than mathematically is not as easy, but let's try.

The lower-left area of Fig. 9-25 shows the intrapulse waveforms to be expected for operation of the TWT high-power amplifier (HPA) with frequency-up-chirped, expanded-pulse RF input. The pulsed cathode voltage is perfectly rectangular with no pulse-top voltage variations of any kind, neither droop nor ripple. The RF input amplitude envelope is also perfectly rectangular. The RF input frequency, however, ramps up linearly from f_1 to f_2 , where $f_2 - f_1 = 500$ MHz, in this example. If the TWT was completely non-dispersive—that is, if phase-delay was perfectly linear as a function of frequency—and there were either a perfectly matched antenna feed or a perfectly matched TWT sever load—it is not necessary for both to be matched in order for there to be no echo distortion—then the relative phase-delay from RF input through the TWT to the antenna feed input would have the smooth parabolic shape shown. (If $f = f_1 + 500 \times 10^6 t/T$, where f is the instantaneous frequency, t is time measured from the pulse leading edge, and T is the pulse duration, then relative phase, where phase is the time integral of frequency, would advance as $500 \times 10^6 t^2/2T$.) But we already know that our system has distortion in the form of the dual reflections from antenna feed and TWT sever load. How will it manifest itself?

The transmitted pulse is now very long with respect to the two-way propagation delay between TWT sever load and antenna-feed input. There will be no time-delineated separate echoes as in the short-pulse case. By the time a signal has traversed the two-way path, however, the instantaneous frequency being

amplified by the TWT will be higher than it was when the reflected signal was being amplified. Therefore, the re-reflected signal, when vectorially summed with the instantaneous primary signal, will appear as a modulation sideband, having an amplitude as before of $0.1 \times 0.1 \times 10 = 0.1$, or 10%, of the "carrier" signal. As this "sideband" rotates around the tip of the "carrier," as illustrated, the resultant signal will experience both phase and amplitude modulation with respect to the carrier. (Note that this is modulation of a carrier signal by a single sideband. This is not the same as a single-sideband broadcast signal, so-called SSB, which transmits only the sideband, or sidebands, with no carrier at all. The complete designation for this kind of broadcast is single-sideband, suppressed-carrier transmission.) The amplitude-modulation index will be 10%. The peak phase deviation in this case will be the angle whose sine is equal to the relative sideband amplitude ($0.1 \times 0.1 \times 10 = 0.1$), which is approximately equal to an angle expressed in radians as 0.1 rad, or 5.7° . The phase modulation or distortion will be cyclic as the phase angle swings from leading to lagging in response to the re-reflected signal vector rotating around the instantaneous carrier, the fixed frame of reference. The modulation rate, or modulating frequency, depends on the pulse duration, T . But the number of phase-modulation cycles is the same regardless of T and is determined only by the frequency change, $f_2 - f_1$, (in this case, 500 MHz) and the time-delay between the two reflections (in this case, 10^{-8} s, or 10 ns), giving a round-trip delay of 2×10 ns, or 20 ns. The number of modulation cycles is the product of the two, 5×10^8 cycles/s times 2×10^{-8} s, or 10 cycles.

Another way of looking at this is to see that the resulting phase and amplitude modulation is the interference pattern, or beat, between a frequency-up-chirp ramp and an identical frequency-up-chirp ramp delayed from the first by 20 ns and having an amplitude of 10% of the first one. (We will ignore the effect of subsequent modulations, each of which will be 10% of the amplitude of the previous one and 20 ns later in time.) The beat frequency will be constant, equal to the amount that the up-ramped frequency would have changed in 20 ns.

This involved discussion is really about an intermediate, almost parametric, set of conditions that affects only the transmitter portion of the system. What actually gets transmitted in the case illustrated are successive long pulses, each with the same frequency up-chirp, each with the same 20-ns delay, and each one being 10% of the amplitude of the pulse that preceded it. If the receiving-system signal processor compresses each frequency-ramped pulse into an equivalent 2-ns pulse, then its output—which is made up of return signals from a single, stationary radar target—will be successive 2-ns-wide pulses that are separated by 20 ns and have 10% (-20 dB) of the amplitude of the previous pulse. Time-delineated pulses don't appear in the composite transmitted signal, but they do appear in the output of the pulse-compressor signal processor.

So far, it has been assumed that the TWT is a linear amplifier, capable of producing an output that is both amplitude and phase modulated as a result of modulation by a single sideband. But suppose, as is the case with virtually all high-power transmitting tubes, that the amplitude is saturated and remains constant, regardless of the apparent change in RF drive power caused by the in-phase (or opposite-phase) component of the sever-load re-reflection. In this case,

the TWT will behave like an amplitude limiter because that's what it is when it is operated to produce maximum power output. How will this added transmission anomaly affect what comes out of the receiver signal processor?

We can gain some helpful background by reviewing Fig. 9-26, which takes a simple look at both amplitude and phase modulation of a carrier signal by a single modulating frequency. To produce pure amplitude modulation (AM) with no residual phase modulation, two equal amplitude-modulation sidebands are required, spaced in frequency above (upper sideband) and below (lower sideband) the carrier by an amount equal to the modulation frequency. The two sidebands can be represented as contra-rotating vectors added to a stationary carrier vector. The upper sideband rotates clockwise. It has a relative angular velocity with respect to the carrier of $2\pi \times f_m$. The lower sideband rotates counterclockwise. It has the same angular velocity but in the opposite direction. Their components that are in quadrature with the carrier are always equal but opposite. Their components that are in-phase with the carrier either add together to the carrier or subtract together from the carrier. Modulation, therefore, is only of amplitude, and the phase angle of the carrier is never changed. In the time domain, the envelope of the transmitted RF signal is seen to have a superimposed sinusoidal modulation. When each sideband is 1/2 the amplitude of the carrier, the modulation index will be 100%. The RF peaks will be twice the carrier amplitude and the valleys will be zero. Each sideband will have 1/4 the




Figure 9-26. Amplitude and phase modulation of an RF carrier by a single modulation frequency.

power of the carrier, the total transmitted power will be 1.5 times that of the carrier, and the peak instantaneous output power will be 4 times that of the carrier.

If the polarity of one of the two sidebands is reversed, however, the situation is entirely different. The two sidebands as vectors will still be contra-rotating with respect to the stationary reference carrier vector, but instead of their quadrature components subtracting from each other, it is their in-phase components that are equal but opposite, and they will cancel each other. Therefore, the predominant result is the modulation of the phase angle of the carrier sinusoidally from maximum lag to maximum lead. The resultant vector of the carrier and the quadrature sideband components is always larger than the carrier, except when the sidebands are opposing one another (one is up and the other is down). This means that there is an unwanted component of amplitude modulation, or residual AM. To reduce it to zero, an infinite set of harmonically related modulation sidebands is required, half above the carrier frequency and half below. To produce either AM or phase modulation (PM), only one sideband is actually required. It does not matter whether the sideband is upper or lower. The purpose of additional sidebands, therefore, is to either cancel completely (AM) or partially (PM) the unwanted modulation type.

Returning to the original discussion of radar echo distortion and how amplitude saturation might affect it, we now can see that suppressing the AM component produced by our single-sideband echo distortion requires that at least a second sideband be generated in contra-rotation with it. The power of this sideband is contributed by the AM that is no longer present. (If the simultaneous AM and PM produced by the single sideband were visualized as simultaneous but superimposed double-sideband effects, as shown in the frequency-domain representation of Fig. 9-26, we can see that the oppositely phased lower sidebands would cancel and the two upper sidebands would add together. If, then, the AM were suppressed, which is what the saturated TWT will do, the two AM sidebands would disappear, leaving only the PM sidebands, whose upper sideband would be only half the amplitude of the former single sideband.)

By introducing an opposite-rotation sideband to our original scenario, we can create an effect similar to antenna reflection that arrives *before* the signal producing it, rather than after it. Now both upper and lower beat frequencies will be present. Their components combine to suppress AM and to enhance PM. But at the output of the receiver signal processor, we now see, instead of echoes that come *only* after the main signal reflection, "echoes" that precede it as well. Even though it appears to defy physics to have radar target returns occurring before the "main bang," as it were, we must remember that there really is no main bang. It is constructed by the signal processor only after the very long frequency-coded surrogate pulse has been completely absorbed by the receiver. Because the expanded pulse can be many thousands of times longer than the compressed "high-resolution" equivalent pulse, there is plenty of anticipatory time in which to synthesize "leading" as well as "lagging" echoes (as described in far more elegant form in "paired-echo" analysis).

So far, all of this discussion has only been a lead-in to the problem of modulator pulse-top ripple. In Fig. 9-27 we see the same TWT, except that there is no

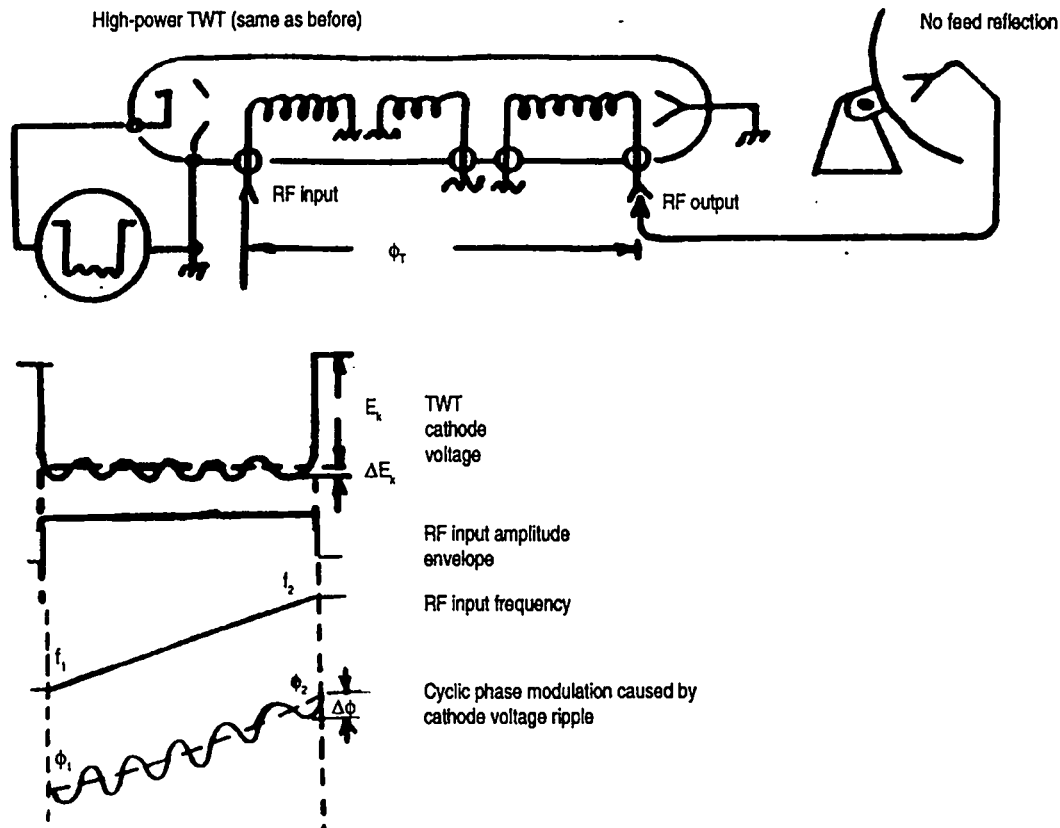


Figure 9-27. Cyclic phase modulation of a TWT through phase pushing.

longer a reflection of RF power from the antenna feed so there is no more phase-angle "pulling." In its place, however, there is now pulse-top ripple from the line-type pulse modulator, which affects the cathode-voltage pulse applied to the TWT. This voltage variation will modulate the velocity of the electrons that produce the TWT beam current, speeding some up and slowing others down. The total phase-delay of the TWT is related to the velocity of these electrons because it is the interchange between the RF fields on the TWT slow-wave circuit and the velocity/density modulation that these fields impart to the electron beam that produces RF gain in the TWT. The length of time that it takes the electron beam to traverse the action space of the TWT can be directly related to phase-delay. The electron velocity, in turn, is related to the TWT cathode, or beam, voltage. The potential energy imparted by the cathode voltage to each electron, which is equal to the product of the voltage and the charge of each electron, is converted to an equal amount of kinetic energy, expressed as $1/2 mv^2$, where m is the electron mass and v is the ultimate electron velocity. Velocity can then be seen to be proportional to the square root of voltage. In other words, a small change in voltage will produce half as much change in velocity.

Were it not for the presence of an RF circuit in the tube that shares in the propagation of RF power from input to output, we could say that the phase-pushing factor of our linear-beam microwave tube—or the amount of per-unit phase change for a given per-unit voltage change—would be $1/2$. A way of stating this more mathematically is that the phase change divided by the total

phase equals $1/2$ the quotient of the voltage change divided by the total voltage. The presence of the RF circuit reduces the phase sensitivity to voltage because that portion of the total RF throughput handled by the fixed-dimension circuit is, of course, independent of electron velocity. Its effect can be accounted for by including a coefficient that can be called a circuit factor. Experimental data shows that such a factor has a value of approximately 0.7, which would give a composite pushing-factor coefficient of 0.5×0.7 , or 0.35 for a typical TWT. A typical electrical length for a high-gain, broad-band TWT is 30 wavelengths, or $30 \times 360^\circ$ ($10,800^\circ$). Therefore, a 1% change in beam voltage can be expected to produce a phase change of $10,800^\circ \times 0.01 \times 0.35$, or approximately 38° , which is pretty much what we would actually measure.

Now that we have gained an appreciation for how cathode-voltage variations will affect, or "push," TWT phase-delay, we can see how pulse-top ripple might affect the frequency-coded, expanded-pulse transmitter. Assume again a system has a 1% peak pulse-top voltage ripple—which is by no means easy to obtain. We already know that the peak phase change will be 38° . This is the phase change that would be produced from two paired, phase-modulation sidebands, each of which has an amplitude of 0.39 with respect to the carrier frequency. Each sideband represents an echo, or "time-sidelobe," that is only 8 dB below the main target return. However, the 1% variation of beam voltage will also produce 2% to 3% power-output modulation, whose effect on time-sidelobe performance is negligible when compared to its effect on the phase modulation. Acceptable "time-sidelobe" performance is typically 30-dB to 40-dB suppression (not all of which is expected to be accomplished by the transmitter). Power-output ripple, therefore, is a big problem inherent in PFN-discharge-type pulsers in a sophisticated radar system.

There are other applications where power-output ripple, rather than phase-modulation ripple, is not tolerable. Free-electron lasers represent one such case. Here, the RF power used to excite the particle-accelerating cavities must have an energy spread that is very small.

While we are on the subject of the characteristics of microwave tubes rather than the transmitter circuits that drive them, let's look at some generalizations concerning amplitude- and phase-pushing factors for three important categories of microwave power tubes: crossed-field devices, klystrons, and TWTs. These can be seen in Fig. 9-28. In the first column is a factor called incremental conductance, which is a normalized factor relating change in current through each device type to change in voltage across it. The crossed-field devices have transfer characteristics similar to biased diodes, or even Zener diodes. As voltage is increased across them starting from zero, current builds up very slowly until the design operating voltage, which is strongly affected by the magnetic field, is approached. Current then builds up rapidly. Around the operating point, the incremental conductance is typically 10. A 1% change in voltage can be expected to produce approximately a 10% change in current. Both klystrons and TWTs have space-charge-limited linear-beam electron guns, in which current is related to the $3/2$ power of voltage. The incremental conductance factors for them is 1.5.

The next column shows incremental efficiency. This is the amount by which dc-to-RF conversion efficiency can be expected to change from a change in oper-

Microwave-Tube Category	Incremental Conductance	Incremental Efficiency Factor	Amplitude-Pushing Factor	Phase-Pushing Factor
Crossed-Field Amplifier	$\frac{\Delta I}{I_0} = 10 \frac{\Delta V}{V_0}$	$\frac{\Delta \eta}{\eta_0} = 0 \frac{\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta \eta}{\eta_0}$ $= 11 \frac{\Delta V}{V_0}$	$\frac{\Delta \phi}{\phi_0} = 0.4^\circ/1\% \Delta I / I_0$ $\Delta \phi = 4^\circ/1\% \Delta V / V_0$
Klystron Amplifier (Linear Beam)	$\frac{\Delta I}{I_0} = 1.5 \frac{\Delta V}{V_0}$	$\frac{\Delta \eta}{\eta_0} = 1.5 \frac{\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta \eta}{\eta_0}$ $= 4 \frac{\Delta V}{V_0}$	$\frac{\Delta \phi}{\phi_0} = \frac{1}{2} \frac{\Delta V}{V_0}$ $\Delta \phi = 10^\circ - 15^\circ/1\% \Delta V / V_0$ (See note 1)
TWT Amplifier (Linear Beam)	$\frac{\Delta I}{I_0} = 1.5 \frac{\Delta V}{V_0}$	$-\frac{0.5\Delta V}{\Delta V_0} < \frac{\Delta \eta}{\eta_0} < +\frac{0.5\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta \eta}{\eta_0}$ $= 2 - 3 \frac{\Delta V}{V_0}$	$\frac{\Delta \phi}{\phi_0} = \frac{1}{3} \frac{\Delta V}{V_0}$ $\Delta \phi = 10^\circ - 15^\circ/1\% \Delta V / V_0$ (See note 2)

Note 1: Typical five-cavity klystron $\Delta\phi/\Delta V = 0.13^\circ/V$, at 8000 V = $10^\circ/1\% \Delta V/V_0$

Note 2: TWT examples: • L-band QKW-1089 0.25°/V, at 10 kV 25°/1% $\Delta V/V_0$

• L-band VTL 5341H 0.2°/V, at 10 kV 20°/1% $\Delta V/V_0$

• Q-band 914H 48°/1% $\Delta V/V_0$

• X-band 8763H 35°/1% $\Delta V/V_0$

Figure 9-28. Amplitude- and phase-pushing factors for different microwave-tube types.

ating voltage. For the crossed-field device, the change is negligible. For klystrons, which tend to be underloaded by design, a 1% increase in voltage can produce a 1.5% increase in efficiency. (This is because the klystron's RF impedance coupled to the output gap from the external load is typically less than that required to develop the maximum permissible RF voltage swing before electron turn-arounds occur at the negative RF voltage peaks.) TWTs are not so easily characterized, but they have experimentally shown changes from 0.5% decrease to 0.5% increase for a 1% voltage change.

The amplitude-pushing factor for each device, shown in the next column, is the sum of the previous factors. For crossed-field tubes, a 1% voltage change, a 10% current change, and a 0% efficiency change total an 11% RF power output change. For klystrons, 1% voltage change plus 1.5% current change plus 1.5% efficiency change gives 4% RF output change. And for TWTs, the comparable sum is 2-3%.

The phase-pushing factors are shown in the last column. Partly because of the crossed-field amplifier's (CFA) initial, relatively low RF gain—which means that the rotating space charge in the circular-format tubes is always strongly influenced by the RF fields produced by the drive power—its phase-pushing factor is relatively low. When it is driven by current, the pushing factor is typically 0.4° for a 1% current change. However, because the incremental conductivity is 10, the factor is 4° for a 1% voltage change. The klystron, unlike the TWT, has no intentional RF field coupling between the cavities distributed along the beam tunnel that make up the RF structure. Phase length, therefore, is directly proportional to electron-beam velocity change. A 1% change in beam voltage will produce a phase change equal to 0.5% of the total phase length of the tube, which is typically 2000° to 3000° , giving pushing factors of 10° to 15° for 1%




Figure 9-29. TWT electrical inputs that affect amplitude- and phase-pushing (see Table 9-1).

voltage change. For TWTs, the phase change for a 1% voltage change will be in the vicinity of 0.33% of the total phase length of the tube. (This calculation includes the RF-circuit factor.) Total phase lengths for TWTs, however, can vary considerably, depending upon gain, bandwidth, power output, operating voltage, etc. The pushing factors, therefore, can be anywhere from 10° to 50° for 1% voltage change. Some examples of actual measured pushing factors are also shown in Fig. 9-28.

Figure 9-29 shows a schematic diagram of a 25-kW X-band, coupled-cavity TWT transmitter. Table 9-1 lists measured pushing factors for all its important electrical inputs. (A large number of these were to have supplied the total RF power for a major phased-array radar transmitter.) As can be seen, the cathode voltage is still the most sensitive input as far as phase change is concerned

Table 9-1. Amplitude- and phase-pushing effects for different electrical TWT inputs.

TWT input parameter	Change in output (dB) for 1% change in input	Change in phase (deg.) for 1% change in input
Cathode voltage	0.1	35
Collector voltage	negligible	0.35
Grid voltage	0.05	5
Filament voltage	Less than 0.1	1.7
Focus-coil current	Less than 0.1	Less than 0.1
RF drive power	0.2	2




Figure 9-30. Measured phase-pushing factors for klystron and crossed-field amplifiers.

because it alone directly affects electron velocity. The second most sensitive input with respect to phase is the grid voltage, which has only a minor influence on electron velocity but plays a major role in determining beam current.

Measurements of klystrons and CFAs are illustrated in Fig. 9-30. For the klystron that has five cavities and a nominal 8-kV beam voltage, a 10% voltage change (800 V) produces a phase change of 100°, or 10° for 1% change. The pushing factor of the X-band CFA is shown to be strongly influenced by the RF drive power, varying from as high as 15° for a 1% voltage change when RF drive is 9 kW, down to a change of 3.8° for a 1% voltage change when the RF drive is 35 kW. These results clearly demonstrate the phase-locking aspect of high RF drive.

What about instances where electron velocity approaches the speed of light? Can relativistic effects alter the performance of conventional high-power RF amplifiers? Relativistic effects can modify phase-pushing factors at voltages that may seem surprisingly modest. A completely general expression for incremental phase change as it is affected by incremental cathode-voltage change includes a relativistic constant, K_R , which can be defined as

$$K_R = \frac{1}{1 + 1.5pV_B + 0.5p^2V_B^2},$$

where V_B is the electron-gun beam voltage and p is the electron charge divided by the product of electron mass and the speed of light squared, or 1.96×10^{-6} . For a tube operating with 10-kV beam voltage, the electrons will travel at less than 20% of the speed of light, and K_R will be an almost negligible 0.97. For a 50-kV beam voltage, the operating voltage for a host of TWT power amplifiers, electrons will still travel at less than half the speed of light, but K_R will have dropped




Figure 9-31. Phase-pushing multiplier plotted against cathode voltage.

to 0.84. For a 100-kV beam voltage—a large number of megawatt-class tubes operate in this regime— K_R is only 0.76. And at 150 kV, where multi-megawatt klystrons operate, K_R is only 0.67. The curve in Fig. 9-31 shows how the basic zero-voltage phase-pushing coefficient of 0.5 diminishes with operating beam voltage.

One solution to pulse-ripple has not yet been mentioned. Faced with an unusably large ripple from a PFN-discharge type pulser, some transmitter designers have tried to further attenuate it by actually clipping off the top the pulse, as shown in Fig. 9-32. For transmitters operating at constant PRF—constant pulse duration is already the rule, determined by the PFN itself—a simple circuit using a clamping diode and a capacitor charge sink can produce significant improvement. An interpulse capacitor discharge bleeder is required to drain the charge from the capacitor so that it has the same beginning voltage for each successive pulse. An added sophistication is a Zener diode reference string, which is only useful at one transmitter operating point. An adjustable-reference dc power supply, whose output voltage is co-ordinated with the pulser output voltage, can make a clipper circuit useful even in variable PRF service. Pulse output power must be sacrificed, of course. It is dissipated in the interpulse interval in the bleed-down resistor. Another drawback is that there will also be an up-tilt or voltage increment to the load-pulse voltage as pulser output current charges up the sink capacitor. Clamp current will be limited by source impedance, which will normally be a characteristic impedance rather than dissipative resistance. If additional impedance must be added in the form of resistance, it

too will contribute to lost pulse output power.

9.3 Modulator output pulse-to-pulse variations

A great deal of attention was devoted to intrapulse voltage variations, which limit a sophisticated radar system's ability to resolve separate targets in range. These variations have the greatest effect on "time sidelobes," or time-domain echoes. Equally, if not more important, is a radar system's ability to discriminate between fixed and moving targets. This ability involves the resolution of "Doppler sidelobes." The efficacy of a moving-target indicator (MTI) radar depends upon the identicalness of successive radar returns from targets whose reflective properties do not change from one pulse to the next. In this way the receiver signal processor can cancel them out. If strong signals from fixed targets can be canceled, then relatively weak signals originating from often far-more-interesting moving targets will not be masked by the "clutter." The "sub-clutter visibility," or "clutter-cancellation performance," of a radar will depend, of course, on the degree to which successive output pulses from the transmitter are identical, regardless of how much intrapulse variation each pulse might have. If there are pulse-to-pulse variations in the transmitter output, returns from fixed targets will have false information imparted to them that might imply target motion.

Any target with a component of velocity in the range direction will reflect




Figure 9-32. Attenuation of pulse-top ripple with a clamping circuit.

signals back to the radar receiver with modulation that is different from the transmitted signal. If the transmitted signals all have an identical wave shape and interpulse interval, even the crudest form of signal processor, which operates only on the detected envelope of the returns, would be able to discern a minute difference between the signals from a fixed target and one with range-rate velocity. If the pulse-canceler delayed successive returns by the transmitted interpulse interval and then subtracted them, there would be a small residue left from the moving-target returns. This is because the apparent interpulse interval of a signal reflected from a target approaching the radar would be a tiny bit shorter than the interpulse interval of the transmitted signal. In other words, there would be a Doppler shift. For the same reason, if the target were going away from the radar, the interpulse interval would be a little longer. In either case, successive pulses delayed by the transmitted interpulse interval would not lie exactly on top of one another and, therefore, could not be perfectly canceled, leaving a residue that could be inferred to mean target range-rate motion. Of course, any pulse-to-pulse differences in transmitter output, such as power output or timing jitter, could also produce a residual output from a pulse-canceler that would suggest false target movement.

Modern, sophisticated radar systems are frequency- and phase-coherent. The transmitted RF is a time sampling of a highly precise, continuously running reference frequency to which the receiver local oscillator and intermediate frequencies are phase-locked. Moving targets Doppler-modulate the RF in the same way that a moving automobile Doppler-modulates the return from a police speed radar. (These are homodyne systems in which the outgoing frequency is mixed with the incoming. The resulting beat, or Doppler frequency, that the highway patrolman measures is proportional to the speed of your car. If too many beats are detected by his radar, the result is your sudden embarrassment.) The MTI radar receiver uses an RF phase detector rather than a simple envelope detector. This strategy, while greatly enhancing the differences between moving and fixed targets, also makes the pulse-to-pulse performance of the transmitter much more critical because pulse-to-pulse changes in the phase length of the transmitter amplifier chain will produce moving-target residue from fixed-target radar returns.

Some radar receivers and signal processors, using range-gating and Doppler "bins," work only in the frequency domain, separating fixed and moving targets by their Doppler signatures. High-performance pulse-Doppler radars operate at very high PRFs, sometimes up to 100,000 pps, so that the PRF spectral lines will be widely separated and Doppler ambiguity will be minimized. (Range ambiguity, on the other hand, is increased in like measure.) This is where the "Doppler-sidelobes" produced by an imperfect transmitter are most directly observable. Power-supply ripple will produce cyclic modulation of the transmitter output, of which the phase-modulation component is likely to be the greatest because the phase-pushing factors of most microwave tubes will produce a greater modulation index than will the amplitude-pushing factor. Unless something is done about it, power-supply ripple will directly affect the voltage to which the PFN of the line-type pulser is charged between successive pulses.

One thing that can be done about this problem is shown in Fig. 9-33. It shows




Figure 9-33. Pulse-to-pulse network charging regulation by deQing.

a PFN with a dc-resonant charging system, as described earlier. The network charge voltage, however, is monitored by a precision voltage divider. The sampled charge voltage is compared with a low-level dc reference signal. When the charge-voltage sample reaches the reference level, an electronic switch, shown as an SCR, is triggered into conduction. This event diverts the charging current from the network and into a dump resistor, which dissipates the residual energy stored in the charging inductor at the end of each charging cycle. This topology is often referred to as a deQing circuit, where the Q that is being snuffed out is that of the charging inductor. However, the dump resistor is wasteful of energy, hence power.

Figure 9-34 shows a more practical circuit that is truly switch-mode, having no intentional lossy components. The charging reactor is given a coupled secondary winding—giving rise to the terms “transactor” or “transductor” for this component—and the power supply must have an output capacitor as charge reservoir. The mechanism is similar to simple deQing, except that the alternative charging-current path is through the secondary winding. Charging occurs when the charge-voltage sample reaches the reference voltage and the SCR switch is triggered. The secondary winding must have somewhat more turns than the primary. The ratio depends upon how great a range of voltage control is required. Note that the regulation point can never be less than half the voltage, because the voltage across the inductor is of the wrong polarity for half of the charge interval. The circuit works best when it regulates close to the full-charge voltage. We will see later how modern high-frequency switch-mode techniques are used to achieve PFN charge regulation.

9.4 Line-type pulser discharge switches

For almost all practical line-type pulsers, the switch will be load-commutated, which means that after the switch has been turned on—that is, when it has made the transition from the initial voltage-blocking, non-conducting state to the high-conduction, low-voltage-drop state—nothing more need be done to the switch. This is because the load current will reach the zero state or even reverse itself of its own accord, even though the switch will still be in the conducting state. This is one of the great advantages of the line-type pulser, because switches of the half-control type (turn-on control only) usually have advantages over their full-control counterparts. In general, the advantages concern cost, complexity, and switching ratio, which is the ratio of blocking voltage to conduction-voltage drop.

Some of the more popular half-control switch types will be now described. The reader will quickly note that there is no mention of sources or current manufacturers of the devices discussed. This is because that at the time of writing (late 1992), no one can assert that any company—even if it is currently engaged in research, development, or production of any of the devices—will be doing so tomorrow. Indeed, much of the discussion relates to past, not present, technology. (Alas, this warning will be all too true of much that remains to be discussed. The golden age of high-power electronic technology is in the past—20 years past, in many instances. What follows will more describe the way things were rather than the way they currently are.)

9.4.1 The hydrogen thyatron

The hydrogen thyatron is a fully enclosed device that is filled with hydrogen gas. It has an anode, one or more grids, and a thermionic cathode. It has been credited with bringing practical electronic control to the timing of radar modulator output pulses. Although “thyatron” and “hydrogen thyatron” are often




Figure 9-34. High-efficiency pulse-to-pulse regulation through charge recycling.




Figure 9-35. Cross-section view of hydrogen thyatron.

used synonymously, thyatron is the more generic term, referring to the early devices that were filled with mercury, argon, and other gases. The charge carriers in a thyatron responsible for anode current are the negative ions formed by the deliberate ionization of the gas that fills it. These negative ions cause conduction from the cathode end to the anode end of the tube, completing the discharge circuit of the line-type pulser.

The advantage of hydrogen thyatrons over other switches is what happens to the positive ions in the plasma. They will be accelerated toward the cathode. The energy imparted to the ions depends upon the conduction-voltage drop between anode and cathode, which is between 100 V and 200 V for a hydrogen tube filled at the proper gas pressure. Cathode damage will not occur until the ions are accelerated through about a 900-V drop. Using mercury vapor or rare gases, the conduction-voltage drop of the tube is sufficient to produce cathode damage through ion bombardment.

Figure 9-35 shows the cross section of a double-gridded, or tetrode, thyatron. It has an indirectly heated thermionic-emission cathode and, usually, a hydrogen reservoir. The electrically heated reservoir is needed because of the high chemical reactivity of hydrogen. (It wants to bond with almost everything, and when it does it is no longer available as a gas to be ionized.) This chemical bonding is called "gas clean-up," and in a high-vacuum electron tube it is a desirable activity. But not in a thyatron. This is because hydrogen thyatrons are designed to be high-power switches. They conduct hundreds, or even thousands, of amperes and, in the pre-conduction state, block tens, or even hundreds, of kilovolts. This is why its grid is spaced so closely to the anode and so far from the cathode. (No, the statement isn't backwards.) The pressure of the hydrogen gas is low, typically 0.5 torr (which is a long way from the vacuumlike 10^{-8} torr or less). Hydro-




Figure 9-36. Increasing thyatron voltage hold-off by means of voltage-gradient grids.

gen, like all gases, has a Paschen's curve describing its voltage hold-off as a function of pressure multiplied by distance. For practical dimensions, operation follows the left-hand side of the Paschen's minimum (see Fig. 7-1). In this regime, at a low, constant gas pressure, spacing must be reduced to increase voltage hold-off. The grid is not only close to the anode, but it is baffled from the cathode and the anode. The holes in the baffle do not line up with the holes in the grid, so there is no direct path for thermionically emitted electrons to travel from cathode to anode. For minimum plasma-arc drop, which is the major determinant of the total conduction-voltage drop across the thyatron, a minimum gas pressure is required. In most large-power tubes, this is accomplished by applying the optimum voltage to the hydrogen-reservoir heater.

Non-conducting-voltage hold-off increases as the spacing between grid and anode is decreased, but there is a practical limit. When the spacing is less than 0.1 in. approximately, another source of charge carriers becomes dominant: electrons. They result from field emission because of the increasing electric-field gradient at tiny anode or grid surface imperfections, especially at the edges of the holes in the grid. (This is the V/R field enhancement discussed earlier.) Field emission usually limits the practical hold-off voltage for a single gap to the 30-kV to 40-kV range.

If a grid-anode gap can hold off only 30 kV or so, how can a designer build a thyatron to hold off 100 kV or more? The answer, which is shown schematically in Fig. 9-36, is to stack a number of gaps in series in the same tube envelope. This is how so-called gradient-grid thyatrons are stacked up to achieve hold-off voltages in the 200-kV to 300-kV range. The voltages applied to the gaps are determined by external voltage division, as shown in Fig. 9-37. (The voltage dividers are sometimes frequency compensated by having an equal-ratio capacitive di-




Figure 9-37. Circuit arrangement using gradient-grid thyatron.

vider in shunt with the resistive divider.) The optimum voltage distribution is not always an equal amount for all gaps, nor is the improvement in total hold-off n times the per-gap voltage, where n is the number of series gaps.

The switching of the thyatron from its voltage-blocking to its high-current-conduction state is initiated by applying a positive-going pulse of voltage to its grid or grids. (In a gradient-grid tube, not all of the gradient-grids are involved. Sometimes none are.) This pulse accelerates thermionic electrons from the heated cathode toward the control grid. On their way, they collide with hydrogen molecules, producing ionization. When the region between cathode and grid has become fully ionized following a multiplicity of such collisions, grid-cathode conduction will occur as a result of the plasma arc, an event described as grid firing. There will be no anode current, however, until the ionized plasma has filled the region between grid and anode, at which time there will be plasma-arc conduction between anode and cathode, which is called anode firing. At that point, the switch circuit will be closed. There is a time interval between grid and anode firing called anode delay-time, which can be as long as a significant fraction of a microsecond, and it is usually not constant. Components of this delay are anode delay-time jitter, which occurs on a pulse-to-pulse basis, and anode delay-time drift, which happens on a longer-term basis. The greater the length of the tube, the longer the anode delay-time and the greater the conduction-voltage drop across the tube are likely to be. This is especially true of multiple-gradient-grid structures. The ionized plasma is characterized by a resistivity just like any conductor, except that its resistivity is inversely proportional to the current through it, tending to make the voltage drop more constant as a function of current than that of a linear resistance. Nevertheless, the plasma resistance will be proportional to plasma length and inversely proportional to its area. The resistivity is a complex function of ion collision energy, mostly between the negative ions (elec-

trons) and the positive ions (protons). The resistivity increases as the hydrogen gas pressure is reduced from its optimum due to gas clean-up, which happens within a pulse as well as over longer time because of the generation of monatomic hydrogen. A single hydrogen atom (H) is even more chemically active than the normal two-atom, or diatomic, hydrogen molecule, H₂. An important difference between the high-vacuum electron tube, which will be discussed later, and the thyratron is that although both use electrons as the negative charge carriers and both have thermionic-emission cathodes, the electron that leaves the cathode of a vacuum tube is almost always the same one that arrives at the anode, whereas in the ionized-plasma device, it is almost never the same one. The electron that does arrive is the last one bumped after a multitude of almost-simultaneous collisions. The higher the plasma current, the less "line-of-sight" the thyratron geometry can be, which is why designers use convoluted cathode shapes that have very high surface-area-to-volume ratios (such as the vane-shaped cathode shown in Fig. 9-35). Another difference is anode delay-time and its annoying by-products: delay-time jitter and drift. These are not present in a high-vacuum electron tube.

To reduce anode delay-time to as low as 50 ns to 100 ns, two control grids are often used, as shown in the tetrode thyratron of Fig. 9-35. Positive-going trigger pulses may be applied to both, but the pulse applied to the grid closest to the anode is sent after the one applied to the first grid by 0.1 μ s or so. The first grid operates with no dc bias, or it can even be positively biased to provide keep-alive ionization current of 50 mA to 100 mA, which is limited by the source. The second grid, however, must have a negative bias applied to it to null any penetration of the electric field below it due to the high voltage applied to the anode. The sequential action of the two grids, or the presence of the keep-alive ionization, greatly speeds the process of total ionization.

When the anode fires, the anode voltage does not fall to near cathode potential instantaneously. During the fall, anode current is building up to its full value. The time integral of the product of anode voltage and anode current represents significant energy per pulse, as much as 25% of the total intrapulse dissipation. Its contribution is directly proportional to pulse repetition rate, regardless of duty factor, and can result in significant average anode-power dissipation even in low-duty-factor applications. This effect is often mitigated by using a saturable reactor in series with the thyratron anode. The reactor allows the anode voltage to drop while it limits the anode current rate-of-rise. Once reactor saturation occurs, after 0.1 μ s or so, the circuit current can build up as before, but now the anode voltage is already down to its steady-state intrapulse value so the time integral of the anode voltage and current product is greatly reduced. This method is another example of a magnetic assist, and it is such a good idea that it is almost continually being re-invented.

Once full ionization has taken place in the thyratron, there is nothing more that the control or auxiliary grids can do. The thyratron will continue to conduct until circuit current is externally terminated or commutated and the voltage across it is momentarily reversed. The ions will then retreat to the surrounding metal surfaces, thus returning the hydrogen to its original, pre-pulse, non-conducting state. Forward voltage can then be reapplied in preparation for the next pulse.




Figure 9-38. Internal construction of a high-power hydrogen thyatron.

Every tube has its own restrictions regarding the required de-ionization time and the rate at which voltage can be reapplied. Repetition rates of several thousand pps, adequate for the majority of radar systems, are readily achievable.

The thyatron makes a good switch. And even though it is not being continually re-invented, it is being continually refined and enlarged-upon. Figure 9-38 shows a sectional view of one of the first major upgrades of the thyatron—the models 3C45, 4C35, and 5C22—which were designed and manufactured by Gen-

eral Electric (who no longer make thyratrons). The GE design replaced the tube's original glass envelope with a metal-ceramic enclosure. The new thyatron had a voltage-enhancing gradient grid, two hydrogen reservoirs, a 200-cm² cathode, and an air- or water-cooled anode. Its anode voltage rating was 40 kV, and its anode current ratings were 2400 A peak, 4 A average, and 75 A RMS. Its average power rating was 100 kW, and its anode-dissipation factor was 55×10^9 . (The anode-dissipation factor, or P_B factor, is derived by multiplying the anode voltage by peak anode current and pulse-repetition rate, which is the rating related to the finite fall time of anode voltage mentioned before.) This tube, which was used to discharge the E-type network shown in Fig. 9-21, drove a 4-MW peak-power klystron transmitter. In this application, the tube's performance was less than its advertised rating. Its peak current was only 1100 A, its average current was 2.4 A, its RMS current was 50 A, and its average power was 24 kW. Its P_B factor was 52×10^9 , however. This was the only performance characteristic close to the tube's rating.

Very much the same geometry was used in a thyatron with even higher performance specifications: the 40-kV, 1-MW-average-power MAPS-40 thyatron. It operated in burst-mode fashion with on-time per burst of 3 to 30 seconds. Its peak current was 40 kA, its average current was 50 A, its RMS current was 1.5 kA, and it had a P_B factor of 400×10^9 . The tube design called for an 8-in.-diameter format, a gradient grid, and both control and auxiliary grids with keep-alive bias. Its cathode area was 5000 cm².

Thyratrons are now being built with cathodes and baffles serving as anodes that can be placed at both ends of the tube and that can be triggered into bidirec-




Figure 9-39. Cross-section view of "back-lit" thyatron.




Figure 9-40. Cross-section view of mercury-vapor ignitron.

tional conduction. In gradient-grid assemblies, these devices allow hold-off voltages of hundreds of kilovolts. Other tubes with "hollow" anode structures can store enough ionized plasma within them to support reverse-polarity conduction for a short time. Another form of hollow-anode structure, shown in Fig. 9-39, is the "back-lit" or "pseudo-spark" thyatron. Filled with low-pressure hydrogen, it, like the conventional thyatron, operates over the portion of Paschen's curve that is to the left of the minimum. The holes in the molybdenum electrodes permit a discharge from the backside of one electrode to the backside of a second, which, because it has a longer path length than a discharge between opposing faces, is the one with the lower breakdown voltage. Ultraviolet light focused through a window at one end of the tube at the backside of the cathode can initiate the long-path breakdown. This breakdown then gives way to what is probably a metallic-ion main discharge between facing surfaces. This discharge can reach tens of thousands of amperes.

Thytrons can also be filled with deuterium (heavy hydrogen), which has twice the mass of hydrogen. At the same gas pressure as regular hydrogen, deuterium will hold off more voltage. And when ionized, its plasma arc is likely to have a lower voltage drop. Deuterium takes longer to de-ionize after current has been interrupted, however.

Another kind of thyatron is the cold-cathode, instant-start thyatron. In these devices, short-term sources of electrons are provided until the thermionic cathode can be self-heated by discharge current. (These devices also use clever

self-energized hydrogen reservoirs.)

Thyratron technology is not dead, nor even dormant. In fact, it has not yet been successfully supplanted by solid-state technology, as we will see.

9.4.2 *The ignitron*

Although thyratrons have been built to support very long pulses (seconds in duration) and charge-transfer per pulse measured in coulombs, the mercury ignitron is often the competitive choice where long-pulse and high charge-transfer are required. Even though ignitron technology is a half-century old, interest in it is periodically reawakened. Modern designs capable of handling peak current in the million-ampere range and charge-transfer per pulse of 1000 C are being investigated.

Figure 9-40 shows the cross-section of a simple ignitron. It comprises a liquid-mercury cathode, an ignitor electrode, and an anode, usually of graphite. The whole device is enclosed in a cylindrical metal envelope that has an insulating anode seal. When voltage is applied to the anode—as much as 50 kV for some types—there will be no conduction. The ignitor electrode, although it is immersed in the liquid mercury, makes no electrical contact because of the surface-tension effect of the mercury. (If it does make electrical contact—in the so-called wetted-ignitor case—the device won't work.) When a pulse of voltage is applied between the ignitor and the mercury pool, an arc is formed, which is called the cathode spot. The cathode spot gives rise to ionized mercury vapor, and a plasma arc between cathode and anode rapidly forms, much as in a thyatron. Once anode current has been established, nothing further done to the ignitor circuit will affect it; the current must be externally terminated or commutated to reopen the switch. When this is done, de-ionization takes place, and the tube returns to its original non-conducting state. (In this respect the ignitron is like the thyatron and all other half-control switches.) To enhance the recovery time of the ignitron, it is customary to heat its anode and cool its envelope. This is done to discourage mercury condensation on the anode and encourage its formation on the envelope. An obvious problem with the ignitron is its lack of deployment versatility; it must be operated vertically so that the mercury pool stays at the bottom and remains in proper relationship to the ignitor. There are, however, ongoing efforts to perfect an "attitude-insensitive" type of ignitron.

Orientation sensitivity is not the ignitron's only problem. It is temperature sensitive in repetition-rated pulse service. Its envelope cooling, which in high-power tubes usually relies on water, requires tight temperature control, especially if the repetition rate increases, pulse conditions are variable, or recovery time must be minimized (or at least stabilized). Temperature affects the mercury vapor pressure, which must be controlled much as the pressure of hydrogen in a thyatron had to be controlled by the hydrogen reservoir. In addition, the ignitron usually demands a more sophisticated internal structure, requiring both control and gradient grids in addition to the basic ignitor of the simple ignitron. Its drive requirements are not trivial either. An ignitron capable of peak-pulse anode currents of 1000 A requires a per-pulse charge-transfer of 10 C or so, and at average currents up to 50 A it might require an ignitor pulse of 500 V (open-circuit) and 100 A (short-circuit) for 20- μ s duration (or 1 J per pulse) in order to

achieve a stable cathode arc spot. In addition, it can require a 2-kV, 3-A control-grid pulse that starts from a negative bias level of 1.5 kV and is slightly time-delayed so that it begins after ignitor current has been well established.

A less obvious problem with the ignitron concerns magnetic induction. It is desirable to use coaxial geometry when mounting the ignitron, as shown in Fig. 9-41. This is not necessarily done to promote faster response. What coaxial mounting does is to keep the mercury-vapor arc on the center line of the tube instead of allowing it to creep up the metal envelope wall. It tends to do this because of the same magnetic pressure that can distort the coils in a PFN. The mercury-vapor arc is a current-carrying conductor, but it is certainly not a very stiff one. By forcing the ignitron's current path to be coaxial from anode to cathode and back to the external circuit, the maximum inductance—which is what the magnetic force is trying to achieve—occurs where the arc channel is on the tube center line. A movement in any other direction would reduce the inductance and is, therefore, opposed by the magnetic force. This same force also tends to shrink the diameter of the discharge and exert an outward pressure on the metal envelope. When peak current is on the order of a million amperes, the magnetic force, if it is not axisymmetric, can easily break off the anode insulator.

Even though ignitron switches require a lot of overhead and are virtually limited to fixed, ground-based applications, they should not be rejected out-of-hand. Where high-energy, high charge-transfer, repetition rates in the kilohertz range, and long-pulse service are required, the ignitron still commands respect for the simple reason that a single ignitron might be able to do the job of a number of less robust individual switches of another type.




Figure 9-41. Coaxial mounting for ignitron switch.

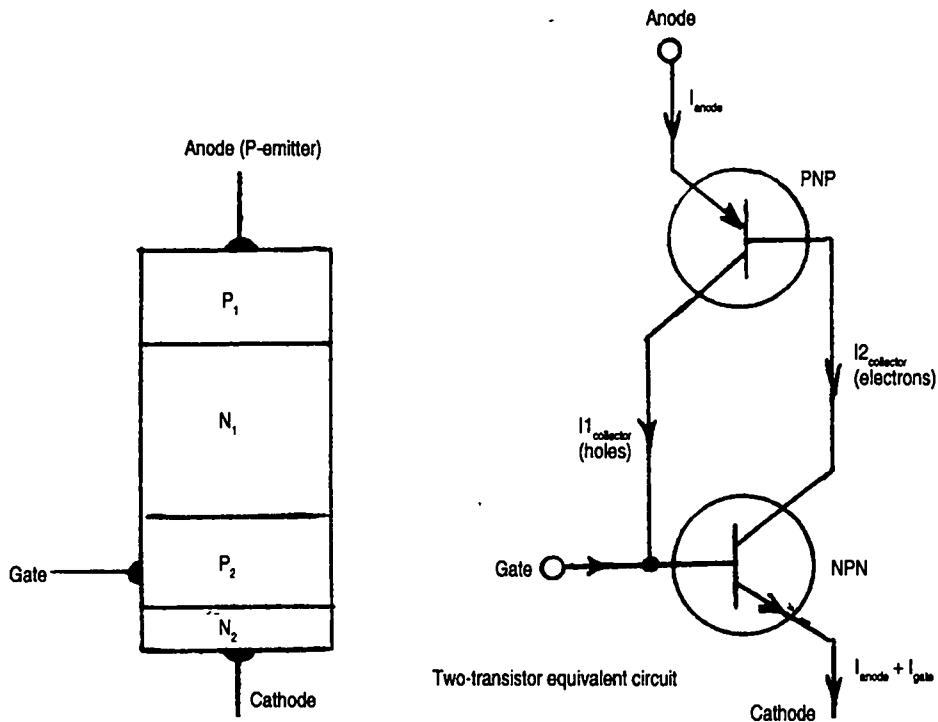


Figure 9-42. The thyristor, or silicon controlled rectifier (SCR).

9.4.3 The thyristor

The thyristor (a term derived from the words "thyatron" and "transistor") is also known as a silicon controlled rectifier (SCR). It would seem to many to be the ideal half-control switch, if only because it is solid-state. It is a four-layer device, as shown in Fig. 9-42, and can be represented as a regenerative cross connection of PNP and NPN bipolar transistors. When applied to the NPN transistor, positive gate current will initiate electron-flow collector current, which just happens to be the base current for the PNP transistor. The emitter of the PNP transistor is the anode for the SCR. Its collector hole-current now regenerates the initial externally supplied gate current, and gate conditions will no longer have an influence on anode-cathode current in the SCR. It is an instant-on device, requires no heater or reservoir power, and, if properly cooled, can have indefinite life expectancy. It has no tolerance for transient overloads, however. Its junction fails only once. (Which is why "indefinite" rather than "virtually unlimited" is the term that should be used when discussing its life expectancy.)

Individual thyristors have been built that can handle kilovolts and kiloamps. What they handle least well is anode-cathode current rate-of-rise following gate turn-on. Unfortunately, this is a characteristic that is important to almost all pulse-modulation applications. What makes thyristors sensitive to current rate-of-rise, or di/dt , is the fact that the gate electrode has a small area with respect to the main junctions. Current flow starts at the gate and must expand radially with finite spreading velocity. Therefore, anode-cathode current begins in a thin, filamentary path near the gate, which uses only a tiny fraction of the total junction area. The spreading velocity has been quantified as approximately 0.2 mm/

μs even when the thyristor has been optimized for it. If current rises too rapidly in the external circuit, it doesn't matter how large the cathode junction is or how much current it can handle under voltage-saturated conditions. The filamentary initial path will overheat and burn out before the first pulse has even peaked. However, the gate can be made with striplike rather than pointlike geometry. Conduction will then begin as narrow strips rather than a small-diameter filament, and the useful conducting area will increase much more rapidly with time. Even so, there is a practical limit to the distance the emitter junction can extend from the strip edge, and that limit is not much more than 0.2 mm if current rise-time is going to be of the order of 1 μs .

(This practical limit to useful area is not fundamentally different from the usefulness of transmission-line conductor thickness, which is familiar to most of us when it is quantified in the frequency domain as skin depth but it is perhaps less familiar when quantified in the time domain as diffusion time. Let's take a coaxial transmission line as an example. At the beginning of a step function of current, all of the current in one direction will be concentrated at the outer surface of the inner conductor and all of the current in the other direction will be concentrated at the inner surface of the outer conductor. As time progresses, current will spread exponentially away from the surfaces, using more and more of the conductor cross-sectional areas. But if the pulse of current is too short with respect to conductor dimensions, there will be conductor area that never sees current flow and is, therefore, wasted).

The thyristor will also exhibit a component of device dissipation not unlike the P_B factor of the thyatron. This is because the voltage drop across it will decrease at a finite rate as the current builds up in the device. Current rise-time will be a function of anode voltage because the transit time of the charge carriers decreases with increasing anode-cathode voltage, although the decrease is by no means linear. This time integral of the products of instantaneous anode voltage and current is aggravated by the spreading time because the dynamic voltage drop of the device will decrease even after the current has almost reached its maximum value. (Typical values in this case for a 1000-A-rated radar-modulator SCR, which was operated at 500 A peak current, showed a 30-V drop 1 μs after current had reached its 10% point, and a 5-V drop 5 μs after the 10% current point.) As in the case of the thyatron, this component of dissipation can be mitigated by an external magnetic assist. Here, the use of a saturable reactor in series with the anode supports the anode-voltage drop before current has built up through the device. The use of magnetic pulse compression has also extended the usefulness of the SCR switch in short-pulse applications. This is because the SCR switches a relatively long pulse, which is sequentially sharpened and magnified by a cascade of saturating reactors, which are often called a magnetic modulator.

Thyristors can be optically triggered as well. Considerable development activity has been devoted to the low-energy (millijoule) optical triggering of kilovolt, kiloamp thyristors by using fiber-optic coupling. This strategy is especially useful for end-of-line power inverters for megavolt dc transmission systems, which require large numbers of thyristors in series. In such an application, trigger-coupling becomes a large problem. Unfortunately, optimizing a thyristor

for low-energy optical triggering while retaining its dv/dt capability following the current pulse—which means desensitizing it to internal convection currents—results in poor current rise-time, or di/dt , which is not what we need in a pulse-modulator switch. Triggering with relatively high-energy laser pulses can result in very high di/dt performance, but the thyristor often becomes trivial within such a laser trigger system. (Its overhead would be greater than an ignitron's.)

So far, we have identified two significant time dependencies relating to turn-on: current rise-time and current-spreading time. There is also a third delay-time, which, again, is not unlike the anode delay-time of the thyatron. This third delay is the time required after gate current has been established, or its photon-induced equivalent, for a sufficient charge to build up within the device bases to support anode current.

The next important time dependency occurs after the pulse current has been externally terminated, or at least brought below the regenerative threshold. This is turn-off time, which is controlled by minority-carrier lifetime and is analogous to the de-ionization time for a plasma-arc device. Recovery time depends upon both the magnitude and duration of the pulsed forward current, and the time is reduced if current is reversed at the end of the pulse. After turn-off, forward voltage will again be blocked so long as its rate-of-change, or dv/dt , is not so great that anode-gate convection current is sufficient to generate a false gate trigger. All thyristors rated for modulator service are characterized according to turn-off time, typically 20 μ s, and dv/dt , typically 100 V/ μ s with the gate open-circuited. (Performance can be improved if the gate-cathode circuit is resistively damped.) The maximum PRF is determined by these after-pulse characteristics.

The upshot of these considerations is that pulse-modulator-optimized thyristors, which can be thought of as dispersion-time-limited—just as RF-rated transistors can be thought of as wavelength-, or skin-depth-limited—tend to have upper, practical per-device physical and electrical limits of 1000 V and 1000 A. This means their maximum power is 1 MW, which translates to 500-kW load power because line-type pulsers are nominally matched-source generators. There are also factors that must be separately evaluated: independent per-pulse energy, charge-transport, and repetition-rate restrictions. Intrapulse transient heating, which is proportional to per-pulse energy, can seriously degrade thyristor life expectancy even in low-duty-factor operation. Operating a thyristor at high PRF, independent of duty factor, can do the same. In either case, the average junction temperature must be kept below 175°C in order to maintain the semblance of reliability. (This implies that peak junction temperature can be considerably greater in pulse service.) For the SCR, the big advantage of line-type pulser service is that the voltage across the switch is minimum when the junction is hottest, at the end of each pulse. The junction will be cooling down as voltage is reapplied.

The overall switching efficiency of almost all half-control devices, thyristors included, is defined as the load power divided by the sum of the load power and total switch dissipation. Efficiency will decrease as pulse duration decreases because turn-on effects will become more dominant. In long-pulse, low-PRF service that reaches saturation junction conduction, thyristor efficiency can reach 98%, about the same as a hydrogen thyatron. (A 50-kV thyatron, switching a

25-kV load voltage, will typically have a 500-V conduction drop, while a 500-V thyristor, switching 250-V load voltage, might have a 5-V conduction drop.) But for short-pulse, high-PRF service, thyristor efficiency can slip below 90% .

Half-megawatt peak-power switching is the niche that single-thyristor switches have settled into, and very few line-type pulsers of lower power use anything but. This is not to say that thyristors cannot be considered in higher-power applications. Indeed, large numbers of individual thyatron chips have been successfully arrayed in hybrid structures to switch tens of megawatts of peak power. Individual thyristor-switched line-type pulser modules have also been deployed to achieve the same result, as we will see.

In today's device nomenclature, the thyristor is also called a reverse-blocking triode thyristor (RBTT), which sets the stage for the discussion of another half-control solid-state switch, the reverse-blocking diode thyristor (RBDT).

9.4.4 *The reverse-blocking diode thyristor*

The thyristor is the family of four-layer devices. The RBDT, originally called a reverse-switching rectifier (RSR), is a device that has no gate electrode, so it is a diode instead of a triode. It is obvious that such a device will not suffer from the gate-related turn-on shortcomings of the RBTT because it has none. But with no gate, how is it ever turned on?

It is turned on by transiently overvolutaging it in the forward direction, which is the direction in which it is already holding off network charge voltage. (The network voltage is typically less than 2/3 of the rated RBDT forward-blocking voltage.) For maximum performance, a large trigger voltage is required. The trigger voltage is added to the network voltage being blocked. (Both the trigger-voltage source and the PFN are isolated from the RBDT by series diodes.) The added trigger voltage forces the instantaneous RBDT anode voltage above its blocking voltage. This trigger voltage must also be carefully controlled to have constant rate-of-rise because this determines the amplitude of the convection current within the device. (The higher the dv/dt , the greater the convection current and the lower the turn-on losses.) In return, however, turn-on occurs simultaneously over virtually the entire junction surface, and the current rate-of-rise is faster for a given RBDT than in any other comparable solid-state switch, except for the high-energy laser-triggered thyristors.

The greater the diameter of the RBDT junction, the greater the minimum differential transient, dv/dt , should be. For instance, for a 0.375-in.-diameter fusion, 5000 V/ μ s rate-of-rise is considered the minimum acceptable. This configuration gives peak discharge currents greater than 500 A at di/dt of 2000 A/ μ s and at pulse durations over 50 μ s. For the larger 0.9-in.-diameter fusion, 15000 V/ μ s is minimum for peak discharge currents up to 5000 A, 2500 A/ μ s rate-of-rise, and at pulse durations up to 20 μ s. Trial-and-error experimentation has shown that slower rates-of-change of trigger voltage lead to premature device failure.

9.4.5 *Spark gaps*

The thyatron, ignitron, and the thyristor family, including RBTT and RBDT, make up the "long-life" family of half-control switches that are suited for continuous, repetitive service. There are other devices comprising air-, gas-, and vacuum-dielectric spark gaps that are capable of even higher products of hold-off

voltage and peak discharge current. For switch voltages above about 300 kV (and into the megavolt range) they are the only choice. They are especially well suited for single-shot, non-repetitive—or at least infrequently repetitive—applications, such as protective crowbar devices. This is because spark gaps, except for vacuum gaps, have inherently longer recovery times—somewhere in the several-millisecond range—making high-repetition-rate operation difficult. Also, all spark gaps suffer from electrode erosion. All devices have an ultimate limit of charge-transfer before they no longer perform as originally intended, and the life of many gap types is expressed in terms of total coulomb-transfer capability.

It is difficult to generalize about spark gaps because most of them are custom designed to meet special requirements. Air, of course, is the cheapest and most readily available gap dielectric. It is self-healing as well. Highly successful air-gap switches are in abundance, especially as single-shot charge diverters, crowbars, or even as low-PRF pulse-modulator discharge switches. (They are especially common in test or experimental modulators where periodic replacement of worn-out electrodes is not a major inconvenience.) Repetitive operation is usually facilitated by blowing air through the gap, often at high pressure, to remove ions and electrode-erosion detritus. This purging also enhances the gap's self-breakdown voltage while minimizing switch inductance. (The power required by a blower varies roughly as the cube of air flow, so this power can quickly become a major design factor.)

Gaps have been built with coaxial and planar geometry. The most common means of triggering the gap is through a "midplane" trigger electrode, usually of the field-enhancement type. These electrodes are pointed or sharp-edged conductors like a rod or thin disc that are mounted so as to be parallel to the equipotential lines that can be drawn between main electrode shapes. With a voltage difference only between the main electrodes, and with the voltage assumed by the trigger electrode established by an external voltage divider so that it is the same as the equipotential surface it shares, there will be no distortion of the electric fields between main electrodes caused by the presence of the trigger electrode. Triggering is accomplished by applying a rapidly rising voltage, with respect to the main electrodes, to the trigger electrode. This will increase the voltage between the midplane and one or the other main electrode and also destroy the equipotential "hiding place" of the pointed end or sharp edge of the trigger electrode, leading to greatly increased electric field at the trigger tip or edge. This, in turn, causes local ionization, or corona, that quickly leads to plasma-arc discharge between the trigger and the overvoltaged main electrode. In most cases where there is already high voltage between the main electrodes (especially if it is a significant fraction of the self-breakdown of the switch), there will be ionization of the region between both main electrodes when the switch closes. (Later we will see an air-gap crowbar that accomplishes main-electrode breakdown even when there is no voltage between them.) This closure occurs within nanoseconds and with virtually no jitter, which is one advantage of the spark-gap switch. The current's rate-of-change is almost always determined by the external circuit.

The triggered vacuum gap (TVG) is a special case. Despite the fact that there is supposedly no conducting mechanism in a vacuum, when an arc does occur in

one it has probably the lowest conducting-voltage drop of any switch considered. (The voltage drop between copper electrodes in a TVG has been measured at less than 20 V for arc currents up to 1000 A.) This is because the charge-carrying ions are derived from the surface vaporization of the cathode electrode. The vaporization is caused by a high-voltage trigger applied to an electrode adjacent to the cathode. (Often this is a rodlike conductor centered in a hole drilled through the cathode electrode.) Delay-time in the TVG, like that of the thyatron, is finite, but it can be reduced by introducing a little puff of hydrogen into the gap as part of the triggering experience. Such an embellishment is called a plasmoid trigger. Titanium hydride can be used to provide the hydrogen puff, which is rapidly gettered after the conduction period. TVGs are more successfully operated at high PRFs than gas gaps. Even PRFs as high as 10,000 pps are not unheard of for TVGs, whereas 100 pps is good going for a gas gap, even with substantial gas blowing.

9.5 Practical system applications of half-control switches

Now let's look at some practical applications of the switches we have just discussed.

9.5.1 The hydrogen thyatron

Although it is presumptuous, considering the rich history of the thyatron, to suggest an archetypal application for it, the high-power pulse modulators used in the Stanford Linear Accelerator Center (SLAC) RF power system are certainly worthy of such consideration. To provide accelerator RF excitation at SLAC, 245 modulators are required, each supplying beam input power to one of the highest-power klystron RF amplifiers ever developed. (It is now being upgraded to even higher power, along with its modulator.) Each pulse modulator, as shown schematically in Fig. 9-43, is elegantly simple. At the time of the original design concept there was no single thyatron switch capable of handling the full power required. However, a single ignitron switch could, and it was given serious consideration. Nevertheless, the thyatron prevailed as the switch of choice,

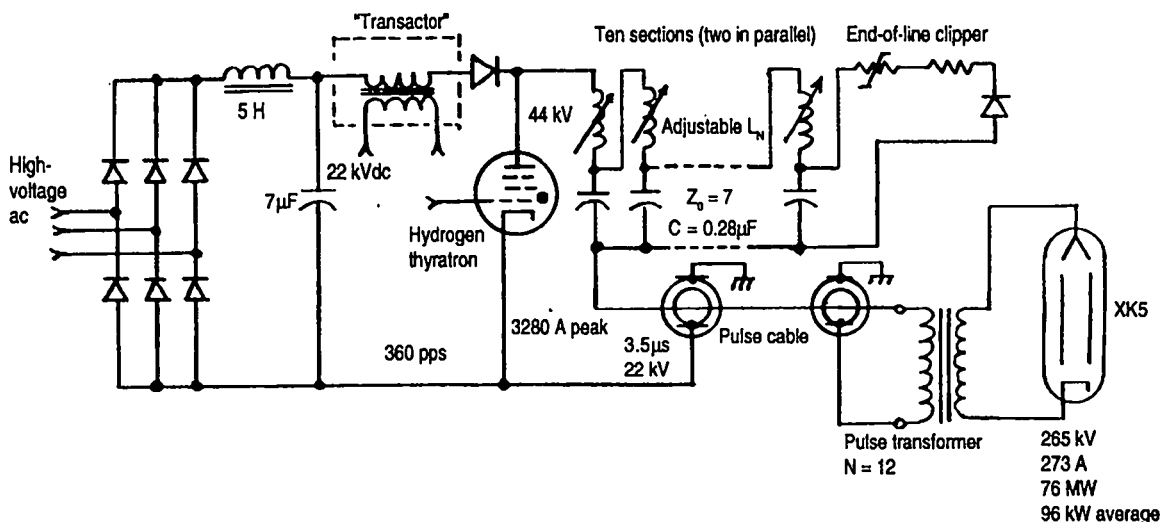


Figure 9-43. Simplified schematic diagram of original Stanford linear accelerator (SLAC) pulse modulator.

even though the design was hedged somewhat to include two parallel 10-section networks. This was done in order to share the load between two individual thyratrons, if two needed to be used. However, thyratrons did become available that were capable of handling the full 41.7-kV anode voltage, 3000-3150-A peak anode current, 3.5- μ s pulse duration, and repetition rates from 60 to 360 pps. In fact, these thyratrons boasted excellent average lifetimes: 70,000 hours at the lowest PRF and even 10,000 hours at the highest.

The parallel-connected pulse-forming networks each comprise 10 identical L-C sections, with the individual network inductors wound solenoidally upward from the capacitor insulating bushings. Each inductor has a copper-slug vernier inductance trimmer. There is no intentional mutual inductance between coils, so the network is in the category of the Rayleigh line. The 10 sections were dictated by the stringent (in the 0.5% category) pulse-top flatness required by RF fields for particle accelerators. The individually adjustable inductance increments affect the pulse shape in accordance with their physical location in the artificial line. (The inductor closest to the front affects the part of the output pulse closest to the beginning.) The PFN impedance is 7 ohms, and the total capacitance is 0.28 μ F. Reverse thyatron voltage needed for forward-blocking recovery is not obtained by network overmatch, as it is in many designs. Instead, the network is undermatched so that the switch is still conducting after the end of the pulse. This allows the output-transformer magnetizing current, which had been building up throughout the pulse, to continue to flow in the network. This current also reverses the network voltage. However, the end-of-line clipper circuit, which tends to oppose PFN voltage reversal, is modified with a series Thyrite non-linear resistance. This resistor has high resistance at the low backswing current but low resistance at high currents, such as those produced by load arcing and load operating-point changes. Regulation of the pulse-to-pulse network voltage is by means of a "scavenging" type deQing circuit that has a "transactor" with a turns ratio of 1.5:1, giving network per-unit charging control from 0.67 to unity.

The original modulator performance was 65.7-MW peak power, 96-kW average power, and 250-kV peak output voltage into the 35-MW XK5 S-band klystron load. Since 1963, the klystron and modulator performances have been continually upgraded. The 1985 second-generation system operates at 152-MW peak power, 136-kW average power, and 350-kV peak output voltage into the 67-MW type 5045 klystron. The thyatron is still the switch of choice, even though anode voltage has increased to 46.7 kV and peak anode current has risen to the 6225-6500-A range (depending on klystron gun perveance).

9.5.2 *The triode thyristor*

If the SLAC application of the hydrogen thyatron is significant, then the thyristor application illustrated in Fig. 9-44 cannot be insignificant, even though it turned out to be only of historical interest. It is the solid-state equivalent of the original SLAC pulse modulator. This application used approximately 1000 thyristors. They were not wired in series, in parallel, or even in series/parallel combinations. They were wired in individual, transformer-combined charging and discharging modules. (This idea was considered unique enough at the time to warrant the issuance of US patent No. 3163783.) The pulse outputs of the 432

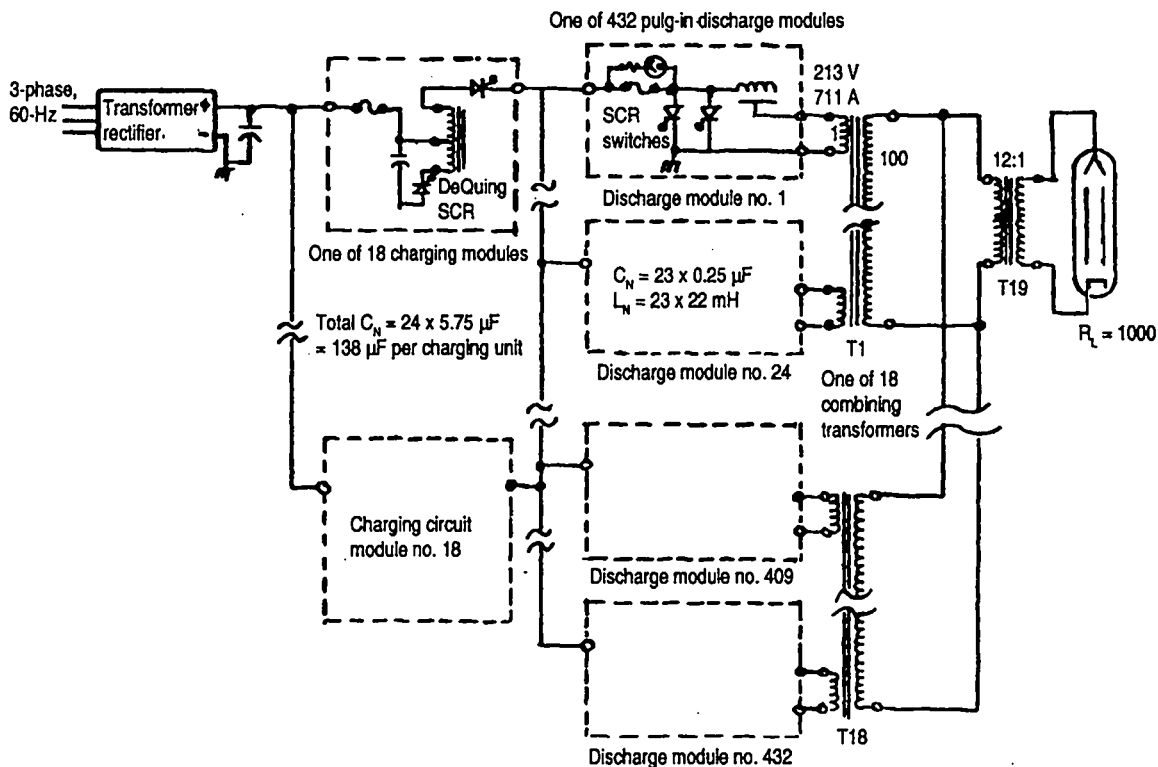


Figure 9-44. Solid-state version of original SLAC pulse modulator.

plug-in discharge modules were combined in groups of 24 by means of 18 combining-pulse transformers, each of which had 24 low-impedance primary transformer windings, at 213 V and 711 A, and one secondary winding. The total primary current was 307,152 A—an impressive number all by itself. The secondary windings of the 18 combiner transformers were all connected in parallel to feed the primary winding of the 1:12 step-up output transformer, which was not so dissimilar to the one used in the thyatron version of the modulator. It had an output voltage of 256 kV at 256-A peak pulse current. (Their product was the 65-MW input to an original SLAC S-band klystron.) The charging-source input to each of the discharge modules was fused so that a faulted module was automatically disconnected from the others. A replacement module could be plugged into the modulator while it was in operation, keeping total modulator availability presumably near 100%, an attribute of the “graceful degradation” exhibited by modular systems of high module count.

The charging of the individual discharge modules was accomplished by 18 individual dc-resonant charging modules that provided thyristor deQuing for charging-voltage regulation and series-thyristor command charging. The triggering of the thyristors was delayed from the ends of the discharge pulses to ensure the recovery of the thyristor from discharge before recharge current was initiated. These charging modules were also fuse protected and replaceable on-line. But the failure of one charging circuit module caused 24 times the impact on modulator performance than that of an individual discharge module. (Note that in any modular system, the effect of the loss of one module on total output power comes down asymptotically to an incremental value of $2/n$, where n is the num-

ber of modules. In the case of 432 modules, one module is 0.23% of the total. But losing it causes the output power to drop by 0.46%, unless the load transformation ratio is also changed to keep output current constant while the output voltage drops.)

Modular systems are not inexpensive, and this one was no exception. Its presumed irresistibility was based on an expected lifetime of more than 10 years for an individual thyristor, compared with a predicted lifetime of 1000 hours for




Figure 9-45. Charging circuit of high-performance solid-state line pulser.




Figure 9-46. Solid-state modulator charging-circuit waveforms.

a thyatron. The long-term operating-cost advantage, it was argued, would make the higher initial-acquisition cost worth it. But the customer didn't buy it.

Another application of the thyristor in a network charging system is illustrated in Fig. 9-45. This is the first reference to a true high-frequency switch-mode power-conditioning application, more of which will be discussed later. Instead of charging the PFN capacitance through a large inductor using only one-half of the L/C resonant period—which is a greatly inefficient use of reactance—the circuit of Fig. 9-45 uses inductance to determine the frequency of discharge of a waveform in which many, many cycles are used in the charging process. The circuit is a half-bridge resonant inverter, where gating $SCR1$ on initiates a full-resonant cycle of conduction through inductor $L1$; the transformer primary; and $C1$. The opposite-polarity half-cycle returns through the anti-parallel diode, $D1$. Conduction of a full cycle of current through the transformer primary in the opposite direction is started by gating on $SCR2$, which also drives $L2$ and $C2$. The waveforms of these two currents, $I1$ and $I2$, and their composite vector sum are shown in Fig. 9-46. The secondary winding of the charging transformer, which is physically small compared with a dc-resonant charging inductor for the

same charging conditions, is connected to a bridge rectifier, the output of which is the unidirectional network charging current. In the system illustrated, part of an all-solid-state modulator, there are four individual such inverter modules feeding the primary of a common charging transformer and rectifier, which, in turn, supply current to five parallel-connected individual discharge modules that will be described later.

The coupling factor of the charging transformer is such that the inverter series-resonant circuits are considerably less than critically damped. The PFN voltage builds up gradually, with each complete cycle of inverter current adding only a small increment of PFN charge voltage. The inverter, however, runs fast. In the system illustrated, the maximum rate is 10 kHz, which all four inverter modules run at for most of the charging interval. As the voltage divider shunting the network announces the fact that the charging voltage is approaching its intended value, the inverter frequency is first slowed to 7.5 kHz, reducing the rate at which charge increments are added to the network capacitance. Then, as the voltage gets closer still, three of the four charging modules are shut down. The network is deliberately overcharged by the last remaining inverter module well before the next discharge event. The network voltage gradually bleeds off. But just before the next pulse, a shunt regulator, which is enabled in time by the interval switch for 600 μ s before the pulse, sucks out one last increment of charge to bring the pre-pulse network voltage to within 0.01% of its intended value. Note that the error-amplifier reference in Fig. 9-45 is not at a fixed voltage but floats with the average value of network voltage. However, it is a time-integrated sample that changes negligibly on a pulse-to-pulse basis, which is all that matters in this application. Although most of the charging system has the attributes of switch-mode regulation, including no intentional power dissipation, the final shunt regulator does dissipate the increment of energy removed by the shunt switch in the resistor, which shunts the charge-coupling capacitor.

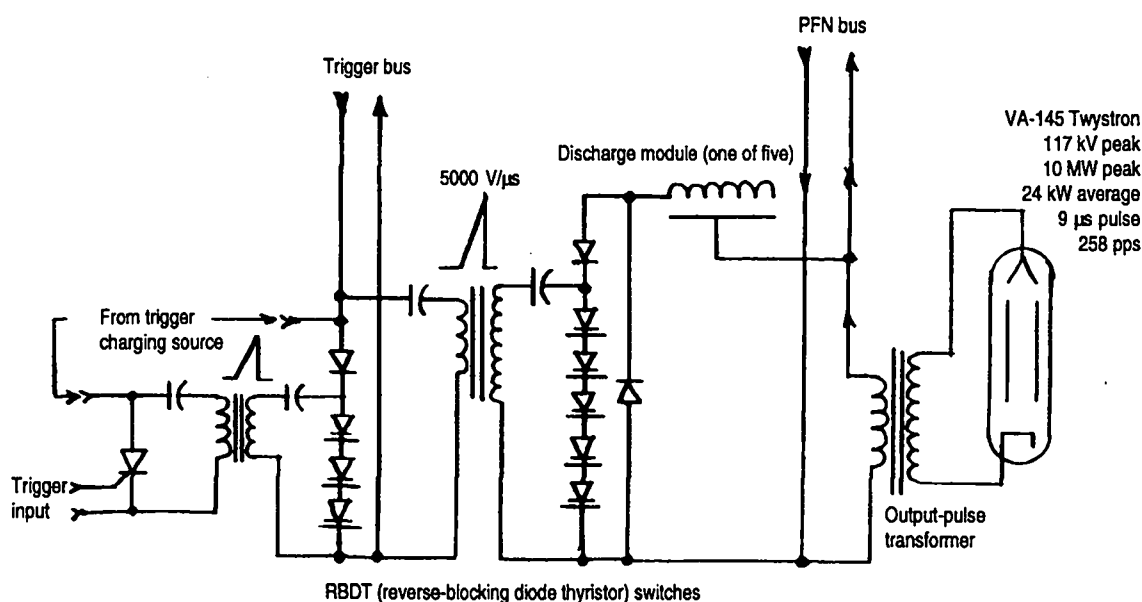


Figure 9-47. Solid-state modulator discharge circuits, using RBDTs as output switches.

9.5.3 The diode thyristor

The solid-state discharge portion of the previously described system is shown in Fig. 9-47. Each of the five parallel-connected discharge modules has a pulse-forming network, a series-string RBDT discharge switch, and a trigger-coupling circuit. A high-voltage, high-ramp-rate trigger voltage is coupled to each module input from a common trigger generator. The trigger generator also uses an RBDT series string as a discharge switch, which is gated in turn by the output of a triode-thyristor (SCR) input stage. The trigger generator networks are pulse capacitors because only the leading edges are of importance. Note that both levels of RBDT switches are isolated from their loads by coupling diodes, so that when the RBDTs are overvoltaged, the coupling diodes are reverse-biased, isolating the loads. But they are forward conducting when the RBDTs avalanche in the forward direction. The combined total output of the discharge modules in this particular system is 10 MW peak. The secondary of the output-pulse transformer provides 1170 kV peak voltage for 9- μ s pulses at 258 pps for an average power of 24 kW.

9.5.4 The ignitron

Illustrated in Fig. 9-48 is an application of the mercury-vapor ignitron as a line-type modulator switch that makes use of its advantages and suffers from none of its disadvantages. This particular system is ground-based and very fixed, so the vertical-orientation requirement of the GL5630 ignitron is no problem. The pulse repetition rate is also fixed at 30 pps, so recovery time is not that important. What is important is the pulse duration of 2 ms, or 2000 μ s. (Needless to say, a PFN having a delay time of 1000 μ s is impressive.) The modulator comprises four vertically arrayed, Guilleman E-Type networks connected in series. The long-pulse capability of this particular ignitron switch, compared with its operating stress levels, is shown in Table 9-2. The charge-transfer per pulse is 2 C, which would be a lot for alternative switch types at the 533-A peak current required. The output energy per pulse is 10 kJ, which is also impressive. (As you may have guessed, this ignitron is no longer manufactured.)

Even more impressive in many ways is the multi-function use of the output pulse transformer. First and foremost, it functions as a step-up high-power pulse transformer, stepping up the 8-kV primary voltage pulse applied to winding N_P , to 100 kV at the secondary winding N_S . An extension of the primary winding, N_X , steps up the primary voltage pulse in autotransformer fashion, providing enough excess voltage to operate a pulse-top clamper circuit (as shown in Fig. 9-32) using clamper diode CR_2 and the power-supply storage fused capacitor bank, C_1 , as the charge sink. A tertiary winding, N_C , serves two more functions. First, it is in series with the network charging circuit. The charging current acts as core-bias current for the transformer, augmenting its volt-time product in order to support the 2-ms output pulse. Secondly, the difference in the winding pulse voltages between N_X and N_C provides the necessary conditions for the inverse, or clipper, circuit, CR_3 and R_1 .

This line-type modulator has been almost completely replaced by the modulating-anode-type pulser, which will be discussed in a later chapter. This modulator is now of mostly historical significance.

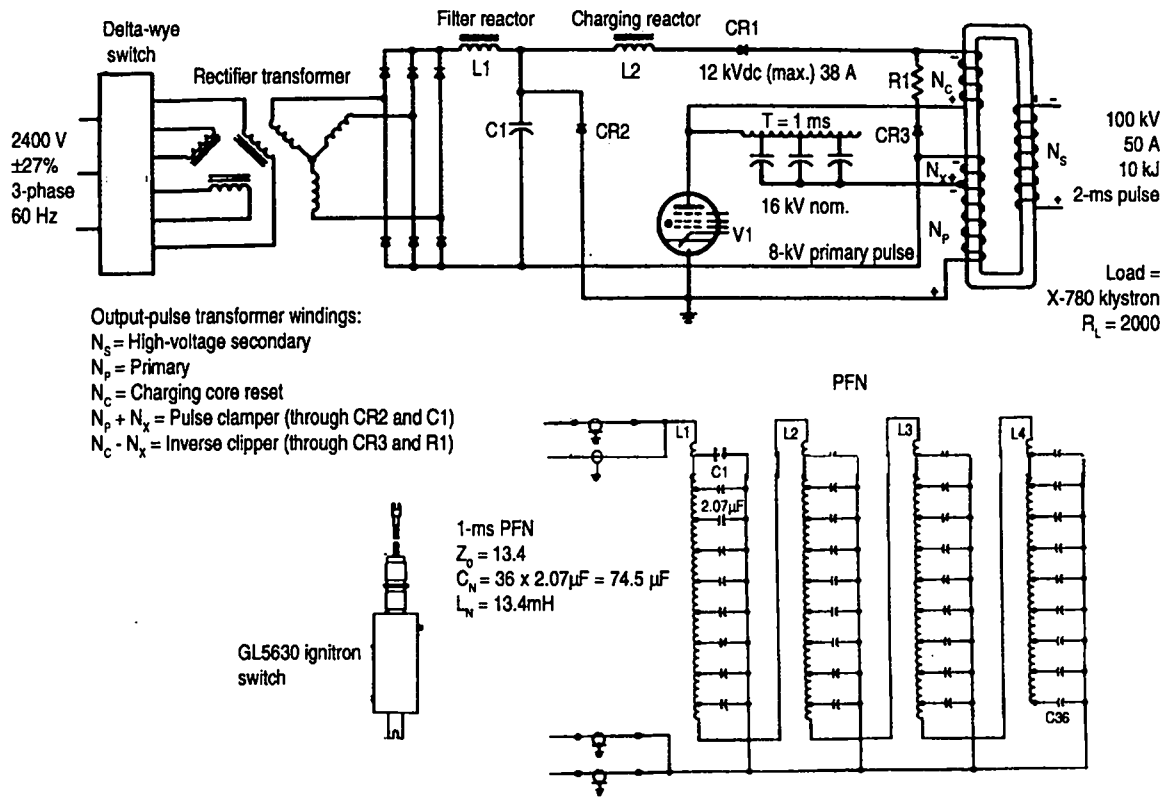


Figure 9-48. Long-pulse line-type modulator using ignitron output switch.

9.5.5 The spark-gap switch

What is probably most surprising about the use of the spark-gap switch in line-type-modulator discharge service is that interest in it is picking up rather than dying down, especially for use in research, test, and experimental applications. A classic example of the use of spark-gap switches—in this case, three of them—is shown in Fig. 9-49. Although capable of high peak power at 180 MW and high energy per pulse at 6.35 kJ, this spark gap's most impressive attribute is average power, 1.2 MW, at 40- μ s pulse duration and 180 pps. The main repetitively pulsed discharge switch, as well as the switch used as a high-speed crowbar, are of the multiple-gap type. Each switch has 20 individual gaps, as illustrated in Fig. 9-50. The multiple-gap switch is not especially low-impedance, having an effective resistance of 2 ohms at the nominal peak-discharge current of 1430 A. For this reason, the impedance and operating voltage of the pulse-forming network are also high: 88 ohms and 252 kV, respectively. The network

Table 9-2. Long pulse capability of GL5630 compared to its stress levels.

	Rating	Service
Anode voltage	25 kV	16.5 kV
Peak current	1000 A	533 A
Average current	50 A	32 A
Peak power	12.5 MW	5 MW
Average power	625 kW	300 kW
Pulse duration	10 ms	2 ms

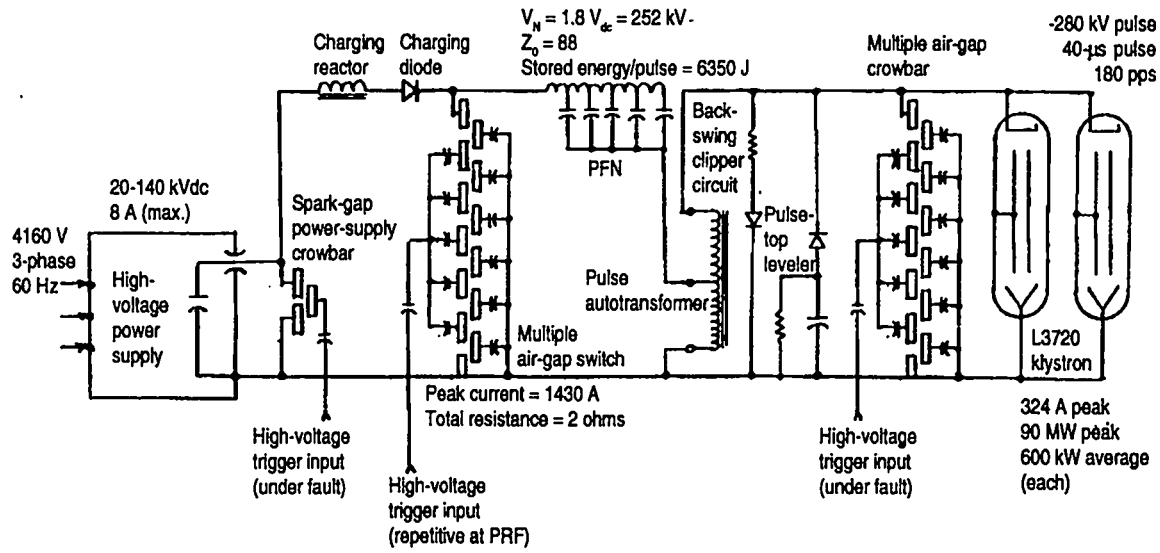


Figure 9-49. Super-power line-type modulator using triggered air-gap output switch.

matched-load output voltage is also high, over 125 kV, so that only modest step-up is required to reach even the 280-kV peak output voltage required by the klystron load. For this reason, an output autotransformer rather than a conventional dual-winding transformer is used. The high impedance levels of the modulator required the use of a smaller, two-gap switch to crowbar the dc output of the power supply. (Note that there is a pulse-top leveler circuit in the form of a clipper diode and R-C self-adjusting charge-sink on the output side.)

The multiple-gap switches are parallel-triggered by a 100-kV peak trigger




Figure 9-50. Details of multiple-stage triggered air-gap output switch.

pulse applied simultaneously to alternately staggered electrodes, as shown in Fig. 9-50, so that there are 20 separate arc channels in series, each having approximately 0.1 ohm arc resistance at the operating current. To facilitate the removal of ion products from the arc channel—a requirement for repetitive operation—an air stream of 8000 ft³/min from a 20-hp blower is funneled through it at high velocity, blowing the arc residue out into the neighboring parking lot. (It should come as no surprise that the use of this modulator was entirely experimental.)

With this we come to the end of our discussion of the line-type modulator and move on the more sophisticated field of the hard-tube modulator.

10. Cathode Pulsers: Hard-Tube Modulators

The hydrogen thyratron, the half-control device most commonly associated with the line-type pulser, is often simply called a "soft-tube" modulator because it is filled with low-pressure hydrogen. If the soft tube is filled with low-pressure hydrogen, what is a "hard" tube filled with? The answer is: as close to nothing as possible. In short, the tube is "filled" with as "hard" a vacuum as can be achieved and maintained, typically 10^{-8} to 10^{-9} torr. (Be assured that the external envelopes of both tube types are equally hard.)

The fewer the number of residual gas molecules in a hard tube the better. A plasma arc in a soft tube is its conduction mechanism. Once the arc is initiated, however, it can be terminated only by means external to the tube, hence its "half-control" designator. A plasma arc in a hard tube, on the other hand, is a fault condition. The conduction mechanism is the flight of free electrons from the cathode, or virtual cathode, to the anode. This mechanism is sometimes called space-charge-limited operation. (The charge carrier in this case is the cloud of electrons that hovers near the thermionic emitter.) Each electron that leaves the cathode must make it to the anode in order to contribute to useful load current. Therefore, any electron collision with residual gas molecules must be minimized.

The number of electrons that are permitted to make the trip per unit of time and per unit of area is called current density. Current density is controlled by means of the voltage applied between the cathode and a helical or meshlike grid (or grids) positioned somewhere between the cathode and the anode. This grid or control electrode of whatever type is placed physically closer to the cathode than to the anode—usually much closer, so that the electric field produced at the surface of the cathode by voltage applied between grid and cathode will be proportionately greater than the field intensity produced by voltage applied between anode and cathode.

10.1 The full-control switch tube

A vacuum tube—or "valve" as the English insist on calling it—that has a cathode, an anode, and a control grid is called a triode. A tube with an additional grid located between the control grid and the anode is called a tetrode. (Other electrodes such as the suppressor grid in the pentode, the shield grid, and even an arc-shield electrode are sometimes employed in hard tubes. But the triode and the tetrode are the basic gridded-tube types.) Their geometries can be planar or coaxial. In the planar configuration, the cathode, grids, and anode take the form of circular discs. In the coaxial configuration, which is also referred to as a radial-beam tube, the electrodes take the form of coaxial cylinders, with the anode being either external or internal (either the outermost or innermost cylinder).

In any case, the total current that leaves the cathode is shared between the anode and the grid(s). The portion of the total current shared with the anode is considered the useful current. Even though grids are perforated so that electrons can pass through them, they are not electronically invisible. If the voltage ap-

plied to a grid is positive with respect to the cathode, some electrons will be intercepted. (This is almost always true for the screen grid and, in a high-power gridded tube when full current is being demanded, for the control grid as well.) One of the goals of tube design is to minimize the control-electrode current because it serves no useful purpose and subtracts from the current that can be useful. The total cathode current in a tetrode can be expressed as:

$$I_{cathode} = I_{anode} + I_{control\ grid} + I_{screen\ grid} = KV^{\frac{3}{2}},$$

where K is the effective perveance. K can also be evaluated as

$$K = \frac{2.33 \times 10^{-6} (A)c}{(S_{diode})^2},$$

where A is the cathode emitting area in cm^2 , c is a constant dependent on tube geometry, S_{diode} is the equivalent diode spacing and given in cm, and V is the effective voltage. The fractional exponent $3/2$ applies to most tube geometries. Effective voltage can be further defined as

$$V = V_{control\ grid} + \frac{V_{screen\ grid}}{\mu_{sg}} + \frac{V_{anode}}{\mu_a},$$

where μ_{sg} and μ_a are the amplification factors of the screen grid and anode, respectively.

The perveance of electron tubes can vary from less than 10^{-6} perv, or 1 μperv (the equivalent to $1 \mu\text{A}/V^{3/2}$), for a high-performance millimeter-wave TWT, to 0.01 perv, or 10,000 μpervs (the equivalent to $0.01 \text{ A}/V^{3/2}$), for a large, gridded power tube. The screen-grid and anode amplification factors, or μ_{sg} and μ_a , are factors that quantify the strengths of the electric-field contributions at the cathode produced by voltages applied to screen grid and anode as they relate to the electric field produced by the control grid. If, for instance, a tetrode had a screen amplification factor of 10, it would require a change of screen-grid voltage 10 times as great as that of the control grid to have the same effect on cathode current change. Similarly, if the anode amplification factor were 100, the change in anode voltage would have to be 100 times as great as the change in control-grid voltage for the same cathode-current change. The screen-grid and anode- μ factors reflect the facts that not only are the screen grid and anode farther away from the cathode than the control grid, but the field produced by voltage on the screen grid must penetrate the shielding effect of the control grid, and the field produced by voltage on the anode must penetrate the shielding effects of both grids. We can see, then, that for positive voltage applied to both screen grid and anode, there will be a negative voltage that, when applied to the control grid, will cancel the field-strength contributions of the others and result in no current leaving the cathode. This balance point is called the cut-off grid bias for a given




Figure 10-1. Progression of soft-tube, line-type modulator into shunt-type, hard-tube modulator.

combination of screen and anode voltages.

Because of cut-off grid bias, the hard tube can be turned off as well as on. Tube current, therefore, can be fully controlled by the voltage applied to the control grid from zero to the maximum. (The maximum current is determined by cathode emissive capability, a factor that varies with cathode area, cathode material, and cathode temperature.) This feature, then, makes the tube a "proportional and reciprocal" full-control modulator switch. It is called proportional and reciprocal because the factors that will reduce current are simply the reverse of the factors that can increase it. More will be said about hard-tube technology later.

10.2 Hard-tube modulator topology

Figure 10-1 shows how the replacement of the half-control switch in the line-type modulator with a full-control switch allows the pulse-forming network to be replaced with a quasi-infinite energy store. This is now possible because the time delay of an artificial transmission line is no longer required to externally turn the switch tube off through commutation. Initiated by a trigger, the pulse duration in the line-type modulator was set by network constants: the PRF was limited by the rate of rise of charging voltage across the soft switch tube, while the pulse-repetition interval (PRI) had to exceed

$$\pi\sqrt{L_{Ch}C_N}.$$

The pulses of current removed from the hard-tube modulator are now replicas of the voltage applied to the grid of the hard switch tube. The output voltage V_0 of the hard-tube modulator is the inverted replica of the circuit's switch tube grid drive voltage. Unlike the line-type modulator, the output-pulse duration and




Figure 10-2. Consequences of pulse voltage decrement, or "droop," on power and phase shift.

PRF of the hard-tube modulator is versatile. On the other hand, the pulse-top ripple of the line-type modulator is replaced by the voltage decrement, ΔV_0 , of the hard-tube modulator, but this condition can be compensated for by ramping the input. Changes in the terminal voltage of the energy-storage capacitor—especially the decrement, or "droop," in voltage that any finite-sized capacitor will experience when a charge is removed from it—can be compensated for by the proportional property of the grid-voltage/plate-current transfer characteristic and the ability of the anode to dissipate significant pulse energy and continuous power.

The capacitance C_S (for storage capacitance) for the hard-tube circuit is many times larger than C_N (for network capacitance) for the soft-tube circuit for a given amount of pulse energy because C_S is only partially discharged each pulse, whereas C_N , as we recall, is completely discharged with each pulse. Even though the loss in capacitor voltage experienced during a pulse can be compensated for by a complementary change in switch-tube anode voltage, as mentioned above, this is an energy- and power-wasteful process. To maximize the efficiency of a hard-tube modulator—which, as we will see, will be difficult enough—the switch tube is typically driven to the minimum anode-voltage drop that is still consistent with the peak current required so that capacitor-voltage droop will be passed on the load voltage. How much droop is tolerable? The answer, of course, depends upon the nature of the microwave tube load. (Refer back to Fig. 9-28 for some idea.)

Figure 10-2 shows the effect of incremental voltage change on pulse energy and phase angle. For a microwave tube having a linear-beam electron gun with space-charge-limited electron flow, beam power will vary as the 5/2 power of beam voltage. For a small value of $\Delta V_0/V_0$, the loss in electron beam power, $\Delta P_B/P_B$, from the beginning to the end of the pulse will be $2.5\Delta V_0/V_0$. The loss in per-unit electron-beam pulse energy will be the time average of the power loss over the pulse, which, for small values, can be treated as linearly varying. This figure amounts to approximately half of the end-of-pulse power loss, or $1.25\Delta V_0/V_0$.

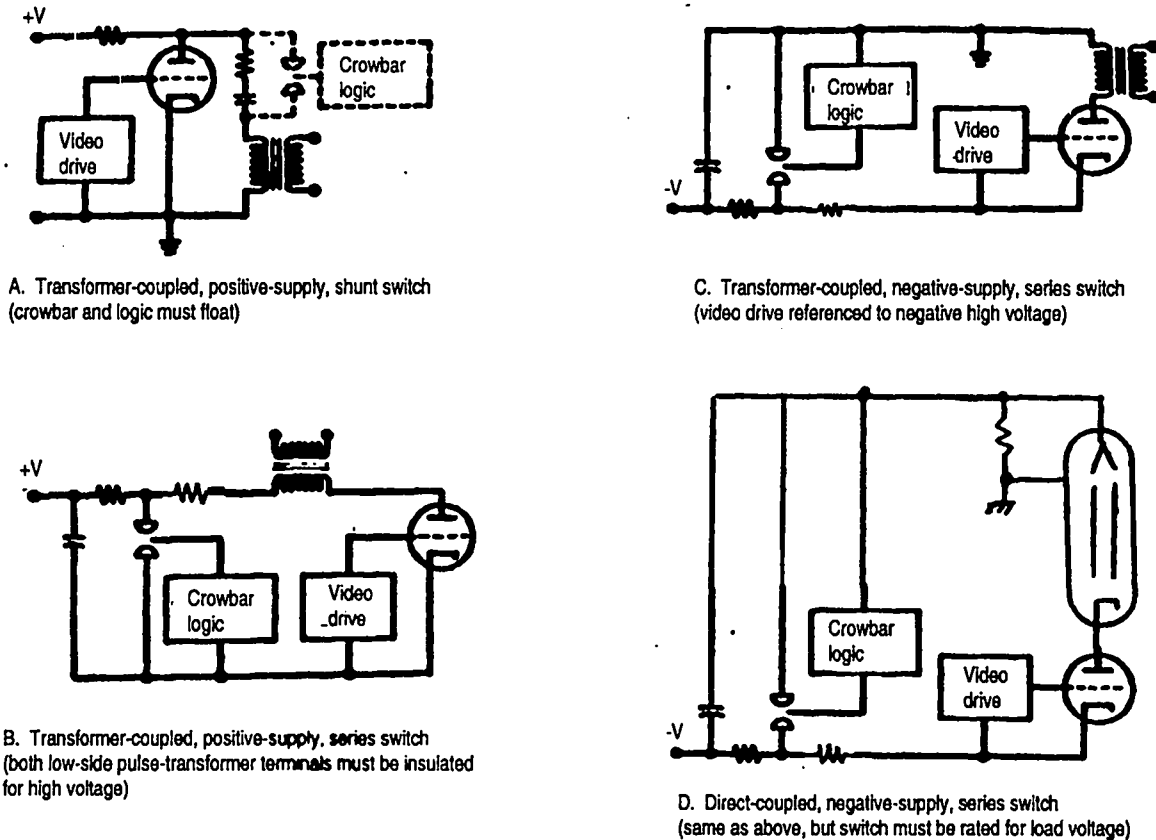


Figure 10-3. Alternative circuit arrangements for hard-tube modulators.

For a crossed-field-type device, the effect of voltage change will be far more pronounced. In this case, $\Delta P_B/P_B$ will be as much as $11\Delta V_0/V_0$, and the approximate per-unit loss in pulse energy will be $5.5\Delta V_0/V_0$. (These values refer only to loss of dc beam.) Other factors listed in Fig. 9-28 relate to how RF power output will be affected.

Because beam-voltage change for small values of $\Delta V/V$ is very nearly linear with time, the change in RF phase angle between input and output of the RF amplifier tube will be nearly linear. (This was noted in the section that discussed phase-pushing effects.) However, linear phase change is indistinguishable from pure frequency translation: the frequency coming out is offset from the one going in. The amount of frequency difference is equal to the percentage of voltage change times the phase-pushing factor expressed in degrees per percentage of voltage change. To find the differential phase change over the duration of the pulse, this product is divided by the pulse duration, or Δt times $360^\circ/\text{cycle}$. In the TWT example shown, the phase-pushing factor is $40^\circ/1\%\Delta V/V$, the voltage droop ΔV is 1%, the pulse duration is 1 ms (0.001 second). The frequency offset, Δf (absolute, not per-unit), will therefore be

$$\frac{1\% \times 40^\circ/1\%}{360 \times 0.001s} = 111 \text{ cyc}/s = 111 \text{ Hz} .$$

The phase slope with time is $40,000^\circ/\text{s}$. Because there are 360° per cycle, the frequency shift, which is exactly the same as a Doppler shift, is $40,000/360$, or 111 Hz. This effect has been referred to by some as the "serrodyne" effect because of its serrated, or sawtooth, waveform, which is caused in this case by voltage removal from a capacitor. Before there were modern frequency synthesizers with phase-locked loops and other means of assuring phase coherency between transmitter and receiver local-oscillator frequencies, the serrodyne effect was actually used to achieve the frequency offset between them. This was done by applying a linear ramp of voltage with its short-duration reset timed so that the period of the resulting sawtooth waveform was the same as the desired receiver intermediate frequency. This sawtooth waveform was applied between helix and cathode of a low-power, continuous-wave TWT that had a relatively high phase-pushing factor. With proper adjustment of the amplitude of the sawtooth waveform, the RF output would be at the desired local-oscillator frequency.

Figure 10-1 showed the most gradual of topological changes that could ease a designer from the line-type, half-control-switch modulator technology into the hard-tube, full-control-switch modulator technology. Figure 10-3 shows four further refinements. The first refinement, illustrated in Fig. 10-3a, simultaneously replaces the soft-tube switch with the hard-tube switch *and* the artificial transmission line PFN with a capacitor energy store, which is coupled with a series fault-current-limiting resistance. In the case of a switch-tube arc, however, all of the charge and energy stored in the capacitor will be dumped. Charge and energy will be transferred to the load connected to the pulse-transformer secondary until the transformer core saturates, which, in a well-engineered application, will not be much longer than the longest normal pulse. Up until this time, the current in the discharge loop will only be about 15% greater than the normal load current, assuming that the normal switch-tube voltage drop was approximately 10% of the initial capacitor voltage and that the load is a $3/2$ -power, space-charge-limited, diode-type electron gun.

When the pulse transformer core saturates, the transfer of energy and charge to the load terminates, but things get much worse for the faulted switch tube. Remember from the discussion on half-control switches that an arc in vacuum is a very good conductor with almost negligible voltage drop. Current is now limited only by the series resistance of the capacitor. What is more, the action integral of the discharge, or

$$\int i^2 dt,$$

is equal to the stored energy divided by the value of series resistance. (joules/ohm is the equivalent of ampere²-seconds.) If the arc in the switch tube was between the anode and one of the small-diameter wires that make up the control grid, the action integral determines whether the grid wire will fuse. If it does and if enough grid wire is involved, the voltage on the grid will no longer effectively cut off anode current in the interpulse interval.

For low intrapulse voltage droop, the capacitor's stored energy should be

many times the amount delivered to the load over the course of the longest normal pulse. ($Q = CV$, so $\Delta V/V = \Delta Q/Q$. For a 10% voltage droop, the stored charge must be 10 times charge delivered per pulse. For a 5% droop, the stored charge must be 20 times charge per pulse. Because the load voltage, V , is almost constant, both stored and load energies are proportional to charge, $W = QV$.) With 10 to 20 times more energy stored in this type of modulator, a means of charge diversion from the faulted switch tube is often necessary. This is usually handled by an electronic crowbar. For physically and electrically large capacitive storage, the topology illustrated in Fig. 10-3a already suffers from the fact that both terminals of the storage system must not only float at high voltage but they must support rapid pulse-voltage change as well. When the added complexity of an electronic crowbar system is factored into this configuration, an alternative architecture becomes attractive.

One advantage of the Fig. 10-3a topology is that the switch-tube grid drive and filament circuits are returned to ground. If we keep this advantage and reconnect the capacitor bank and crowbar so that they are also referenced to ground, the topology of Fig. 10-3b emerges. Although the storage terminals no longer float, both terminals of the pulse-transformer primary float. This means that during the interpulse interval both terminals must be insulated from the core for the power-supply voltage. This introduces the possibility of a failure mode in a component that should normally be one of the most reliable in the entire transmitter. Therefore, this is not a popular solution among those who make their livings designing and building transformers. For this reason, the arrangement shown in Fig. 10-3c is the most common.

In Fig. 10-3c the pulse-transformer primary and capacitor-bank/crowbar are referenced to ground, and the power-supply polarity is reversed. The switch-tube drive and filament circuits are now referenced to the negative high voltage. The filament power supply for the switch tube, which is often no more than a step-down power transformer (albeit one that may be rated for kilowatts of heater power), must now also be insulated for the operating anode-cathode voltage of the switch tube. This is not, however, a high-performance requirement, in that leakage inductance and interwinding capacitance of the transformer are of little or no concern. What might, under other circumstances, be considered excessive leakage inductance can even be advantageous in this case, in that it will automatically limit cold-filament in-rush current. (The resistance of a switch-tube filament, like that of most wires, follows a positive temperature-coefficient characteristic, where resistance increases with temperature—except that the temperature of this wire will eventually reach about 1000°C.)

The grid-drive circuits also will require primary power that is isolated from ground for the full anode-cathode operating voltage. This power is often derived from the isolated filament power (or vice versa). More important, the low-level gate, or on/off trigger signals, must also be coupled from ground-level reference to high-voltage deck reference. Modern designs rely almost exclusively on optically coupled links for this function, even though the common-mode rejection requirements are far less for this application than for the "floating-deck" type modulator, which will be discussed later. In this case, even though the reference deck is at high voltage with respect to ground, this voltage is relatively constant

with time. Any changes consist mostly of the capacitor-bank droop and the voltage drop across the surge-current-limiting resistors caused by the normal peak-load current. Before there were dependable high-performance optically coupled signal links, special pulse-transformer designs and even pulsed RF links were used to transmit low-level timing signals across the high-voltage gap, examples of which we will see later. As complicated as these circuits are, they are in most cases preferable to solving the problem of insulating the pulse-transformer primary winding for the full switch-tube operating voltage—especially because the performance of the pulse transformer is either the most important issue or second only to that of the switch tube itself.

Circuit complications are also unavoidable in the final configuration, which is shown in Fig. 10-3d. This is the direct-drive hard-tube modulator—the top of the line, as it were. In this design, there is no pulse transformer at all. This feature eliminates the limitations on pulse duration, rise-and-fall times, intrapulse droop, and leading-edge overshoot and ringing. The price paid for this improvement is increased circuit complexity and the fact that the switch tube itself must handle the full operating voltage and current of the microwave tube load as separate and independent ratings—not as a peak-power product, where lack of voltage hold-off can be compensated for by increased current-handling (or vice-versa) with any mismatch handled by an appropriate pulse-transformer turns ratio. This direct-drive connection is capable of the highest performance of all modulator types. Given enough excess capability, or head-room, it can be programmed to compensate for almost every shortcoming of the rest of a transmitter, including power-supply ripple, noise, and capacitor-bank droop. In fact, there is a class of switch tube, which will be described later, that will do much of this with no programming at all. Not surprisingly, therefore, this class of modu-




Figure 10-4. The back-swing clipper circuit as used with transformer-coupled unidirectional loads.

lator is almost always the most expensive, if not always the largest.

Before leaving the high-power output-pulse transformer behind, it would be negligent not to discuss its use in a hard-tube modulator that has a diode-type, or unidirectional, conduction (dc) load. We have already mentioned what can be done with energy stored in the transformer magnetizing inductance when used in a line-type modulator. So long as the switch tube is still in the forward-conduction state, the energy will manifest itself as reverse voltage on the pulse-forming network. This is not possible in a hard-tube modulator, because the means of terminating a pulse is to stop switch-tube conduction. This situation leaves no place for the flow of current, which has built up in the magnetizing inductance over the course of the pulse. For this reason, a back-swing clipper circuit, as shown in Fig. 10-4, is required. Without it, the voltage across the transformer primary will jump up at the end of the pulse in a direction that increases the forward voltage across the switch tube (just as it is trying to recover its voltage-blocking state and after its electrodes have been heated to their maximum transient temperatures). The result will be a switch-tube arc. The clipper circuit provides a path for the magnetizing current after the switch tube is cut off. This clipper current can be defined as

$$i_{clipper} = \int I_M e^{-tR_C/L_M} = \frac{I_M L_M}{R_C},$$

where R_C is the clipper resistance, L_M is the transformer magnetizing inductance, and I_M is the peak value of magnetizing current that occurs at the end of the pulse of voltage applied to the transformer. The energy stored in the transformer magnetizing inductance, W_M , can be defined as

$$W_M = \frac{1}{2} L_M I_M^2 = \frac{1}{2} L_M \left(\frac{V_p \tau}{L_M} \right)^2 = \frac{V_p^2 \tau^2}{2L_M},$$

where V_p is the voltage pulse output of the switch tube and τ is the pulse length.

There will still be a transient overvoltage due to the clipper resistance in series with the clipper diode. Without series resistance, the clipper time-constant could theoretically be infinite because it is L_M/R_C . The greater the value of R_C , the shorter the recovery time will be but the greater will be the peak reverse voltage. The volt-time products across a transformer winding must be equal in the positive-going and negative-going polarities. If they are not, the result is time-averaged voltage across the winding—which ultimately is just a length of wire—and it is difficult to maintain average voltage across it. The larger the value of clipper resistance, the shorter the decay time of magnetizing current but the greater the transient overvoltage, which is equal to $I_M \times R_C$. There can be an advantage to using a non-linear resistance such as Thyrite, or the more modern metal-oxide varistor (MOV), as the clipper load. A properly selected non-linear resistance will have nearly constant voltage across it throughout the interpulse interval, thus minimizing the stress on the switch tube immediately following the

pulse, which is the time when the tube is least able to handle it. The MOV clamp voltage should be chosen so that the product of it and the minimum interpulse interval is as close to the product of pulse voltage and pulse duration as possible, making positive- and negative-going volt-time products equal across the transformer winding.

10.3 Storage capacitors

We all know that a capacitor is a device that consists of no more than two plates separated by dielectric material and is capable of storing energy in the electrostatic field between the plates. Nevertheless, entire books have been written on capacitor technology. What a transmitter designer is most concerned with is what the ratings of a capacitor mean and how the stress levels imposed affect capacitor performance.

Generally speaking, the factor most important in determining a capacitor's volume and cost is the energy it is capable of storing. This assumes, however, that once a capacitor is charged, its repetitive discharges will only be partial. (This was clearly not the case with the capacitors used to store energy in the PFNs of the line-type modulators. These capacitors were completely discharged each pulse during normal operation.) In addition, under some fault conditions the voltage can be reversed across the capacitor, which is one of the most stressful events a capacitor can experience—and this applies to capacitors that are *not* inherently polarized, such as electrolytic types. These caveats both hint at the fact that even though the energy stored in a hard-tube modulator capacitor bank might be 10 to 20 times that stored in the PFN capacitance for comparable pulse energy, we are not talking about the same kind of capacitors. Pulse-forming-network capacitors are high-stress components whose failure rates cannot be overlooked. (Water-cooling such capacitors is not unheard of, and most of them are more correctly rated in shot life or number of pulses in a lifetime rather than number of operating hours in a lifetime.)

On the other hand, capacitors for partial-discharge, or dc-filter, service are usually rated for a given operational life, which can be tens of thousands of hours at a given terminal voltage. They come in standard case sizes that differ in energy-storage capability. The height of their bushings also varies, depending on operating voltage independent of stored energy. These capacitors are expected to be operated at their rated voltages. The deratings for life expectancy, which primarily affect voltage stress in the dielectric material, have already been designed in. Operating them at reduced voltage will, of course, result in great extension of life but at the expense of their stored energy, which varies as the square of the voltage.

However, these capacitors are no more tolerant of voltage reversal than any other kind. They will be fully discharged by a short-circuit load fault or by the operation of an electronic crowbar switch. It is extremely important that there be sufficient resistance in series with the discharge loop so that the discharge is critically, or overly, damped in order to preclude voltage reversal. This means that the series inductance of the loop cannot be ignored. In most cases, the value of series resistance will depend upon the permissible peak fault current and will be safely above the value required for critical damping,

$$R = 2\sqrt{\frac{L}{C}}$$

But it is always prudent to calculate, or measure, the actual value of discharge-loop series inductance just to be sure. (This includes the internal series inductance of the capacitors themselves, which can vary depending upon how they are internally constructed.)

In at least one case, an energy-storage system with a greatly undermatched artificial transmission line was used in lieu of a low-inductance capacitor bank having the same value of capacitance. In this design, characteristic impedance was 10% of the load impedance. The delay-time of the network was made equal to 1/2 of the longest pulse duration. The load—a super-power, long-pulse klystron—was connected in series with a hard-tube modulator switch, much like the one shown in Fig. 10-3d. When the modulator switch was gated on to initiate a pulse of load current, there was an instantaneous drop in network voltage of approximately $0.1V$, where V is the initial voltage across the network. The voltage across the klystron and modulator-switch combination was $0.9V$. This voltage would persist for a time of $2T$, where T is the delay-time of the network. At $2T$, the voltage would drop again by an additional $0.1V$, to $0.8V$. If, however, the modulator switch tube was turned off after time T but before time $2T$ had elapsed, most of the network would have been discharged by an amount $0.2V$, but not all of it.

The load current was $0.9 V/R$, where R is the effective resistance of the klystron in series with the conducting-state switch tube. If the storage system used the same value of capacitance but only the minimum-achievable value of inductance, the terminal voltage of the bank would start to droop in essentially linear fashion from the beginning of the current pulse to the end. If the starting voltage is V and the initial load current was V/R , the total droop would equal to the charge removed divided by the total capacitance, or $\Delta V = I\Delta t/C$. (This discussion is simplified by assuming that current is constant throughout the pulse, which it isn't.) The time increment, dt , is the pulse length, $2T$, which was the same as \sqrt{LC} . The total load resistance was $10Z_0$, where Z_0 was the characteristic impedance of the network described above, or $\sqrt{L/C}$. So $R = 10\sqrt{L/C}$ and $\sqrt{L/C} = 0.1R$. Combining things,

$$\Delta V = \frac{I\Delta t}{C} = \frac{V}{R} \times 2 \frac{\sqrt{LC}}{C} = \frac{V}{R} \times 2 \sqrt{\frac{L}{C}} = \frac{V}{R} \times 2 \times 0.1R = 0.2V,$$

which was the same as the end-of-pulse step from the distributed network described above. The charge and energy removed from the capacitance was the same in both cases, except the finite- Z_0 network voltage changes in stair-step fashion while the minimum- Z_0 capacitor bank changes in continuous ramp fashion.

The purpose for creating this artificial network was to minimize the anode dissipation of the modulator switch because the source voltage was constant throughout the duration of the pulse, at $0.9V$. In order to produce a flat-top load

voltage pulse in the latter case, the ramplike, or droop, component would have to be subtracted from the output. To keep klystron voltage constant at $0.8V$, the switch-tube anode voltage would vary from $0.2V$ at the beginning of the pulse to a theoretical minimum of zero (assuming no head-room requirement). The additional anode dissipation per pulse would be $0.1V \times I$, where $0.1V$ is the time-averaged anode drop and I was the load current, as before.

This may have sounded like a good idea, but it wasn't. Its fatal flaw was that it didn't account for what happens under the major fault condition: simultaneous klystron and switch-tube arc-breakdown, or load shoot-through. The only good news was that the fault current was limited by the network characteristic impedance to 10 times I , where I is the normal peak load current. Without series resistance, however, the network stored energy would tend to oscillate back and forth, completely reversing the voltage on the network capacitance periodically until the load arc was finally extinguished. To completely absorb network energy without capacitor-voltage reversal, a series resistance equal to Z_0 , or $0.1R$, was required between network and load. But this completely subverted the reason for the network in the first place. (This is but another example of having to design for the worst-case scenario rather than the hoped-for "normal" operating conditions.)

Rarely, if ever, are the capacitive energy-storage requirements of even a modest-sized transmitter met by a single capacitor (or "can," as it is sometimes called). Multiple cans connected in parallel or series-parallel are the rule rather than the exception. The danger in multiple-capacitor arrays is what can happen if a single capacitor internally short circuits. If the total surge-current-limiting resistance is distributed throughout the array by having at least one resistor capable of dissipating the energy stored in a capacitor in series with each capacitor, a short-circuit failure will cause no damage because virtually no energy will be dissipated in the faulted capacitor. Such an arrangement is so fault tolerant that the failed capacitor may not immediately reveal itself. (If the capacitors are all in parallel, the resistor in series with the failed capacitor will be connected across the dc power supply when an attempt is made to turn it on again. What is left of that resistor will usually be an adequate tell-tale sign for where the trouble lies.)

Series-parallel connections are often employed when the system's operating voltage is much above 50 kV. This is because energy storage density tends to degrade at much higher voltage, and more volumetrically efficient banks can be assembled from capacitors rated at less than 50 kV—by no means a hard and fast number—than those rated at, say, 100 kV. (In fact, probably the most volumetrically efficient energy storage bank assembled for 100kV-plus operation used a multitude of electrolytic capacitors that were individually rated at 450 Vdc. They were assembled in over 200 series-connected tiers and protected from individual overvoltage by a string of Zener diodes shunting them.) When there is a large number of series-connected tiers—say, at least five—the presence of a shorted capacitor in one tier will be far from obvious, even though it will short-circuit all of the parallel-connected capacitors in that tier through its series resistor. This short will cause the overall bank capacitance to increase, but all of the remaining capacitors will be subjected to 25% overvoltage, or less, at the operating voltage. For this reason, some capacitor banks are interconnected with fusible links in-




Figure 10-5. Typical series/parallel capacitor bank (two groups of 25 parallel-connected capacitors) totalling 70 μ F, with triggered air-gap crowbar switches in foreground.

stead of individual series resistors. Fusible links have the advantage of automatically disconnecting a failed capacitor. The fuse, however, must be of the type that permits very little let-through energy, meaning that it must be arc-quenching. A capacitor subjected to too much let-through energy is likely to see its can rupture or its bushing explosively expelled—a spectacular but very hazardous event. When a capacitor bank has only two tiers, the effect of an individual shorted capacitor is the most easily observable, but the penalty for not observing it is the greatest: the possibility of a 100% overvoltage of the remaining tier. The voltage rating of a capacitor is associated with life expectancy more than external arc-over, and many high-voltage capacitors will endure 100% overvoltage for a short time. With only two or three tiers to the bank, it is advisable to measure voltage balance between tiers continuously and provide this information to the high-voltage dc interlock system of the transmitter.

Stacked, multi-tiered capacitor banks also require corona shielding. A capacitor designed to operate at 50 kV may (or may not) have can and bushing geometries shaped so as to be corona-free (see Chapter 7). The capacitor cans of the first tier will typically be connected to the high-voltage return bus, the high-voltage-system “neutral,” which will be at or very near ground potential. The cans of the capacitors in the next tier will be at 50 kV with respect to ground, and their high-voltage-bushing terminals will be at 100 kV with respect to ground,

and so on up the stack. The radii of curvature that are required for corona-free operation continually increase, but the capacitor can and bushing geometries do not, because most practical banks are built up of identical capacitors. Capacitors, therefore, are usually mounted on "rafts" that are surrounded by oblong, donut-shaped corona shields whose cross-sectional diameters increase as the rafts ascend in height and voltage. An example of such construction is shown in Fig. 10-5. Please note that if spacings between conductors are adequate, very little shielding is required to prevent voltage breakdown. This is because corona under these conditions is self-healing; the exterior of the ionization sheath tends to expand so as to have the same effect as a large radius-of-curvature surface. However, in today's sophisticated systems, which involve ever-more-sensitive and susceptible low-level circuitry, the electrical noise resulting from the corona dis-




Figure 10-6. Capacitor charging energy relationships as a function of amount of capacitor discharge.

charge is increasingly intolerable. (Remember what we used to pick up on the AM car radio when we drove near high-tension wires, especially on a foggy day?)

We have discussed the fact that the removal of charge from a capacitor bank, no matter how large, will result in its terminal voltage being smaller than it was before the charge-removal interval, or pulse. In the interpulse interval, therefore, it will be necessary to replace the charge removed during the pulse. In the discussion of line-type modulators, much was made of the issue of recharging without incurring energy loss if the current was limited by resistance. Is there a similar concern when it comes to replacing partial discharge? The answer is not nearly as much. The reason for this is shown graphically in Fig. 10-6, which plots energy removed from the dc power supply, energy transferred to the capacitor bank, and energy dissipated in series resistance versus the per-unit amount that the capacitor is discharged each pulse. (This graph assumes that the resistance is small enough so that the RC time-constant is less than approximately $1/3$ of the minimum interpulse interval, assuring that the capacitor will recharge to approximately the power-supply terminal voltage between pulses.) If we look at full discharge, or per-unit discharge of unity, the results are familiar. The per-unit capacitor energy transfer and resistor energy dissipation are both unity, and the per-unit energy removed from the power supply is 2. Notice, however, that only the energy removed from the power supply is a linear function of per-unit discharge voltage. For less-than-unity per-unit discharge, energy transferred to the capacitor falls off less rapidly than dissipation in the resistor. Therefore, the efficiency of energy transfer increases as $\Delta V/V$ decreases.

The reason for this can be seen if we note the change in charge during a pulse of current, $\Delta Q = C\Delta V$, and the change in capacitor energy, $\Delta W = \Delta Q$, times average capacitor voltage during discharge, which can be expressed as

$$C\Delta V \times \frac{V + V - \Delta V}{2}$$

or

$$CV\Delta V - \frac{C(\Delta V)^2}{2}.$$

The energy removed from the power supply during recharge is ΔQ times the power supply voltage, which can be expressed as $C\Delta V \times V$. Recharging efficiency is the ratio of energy transferred to the capacitor to the energy removed from the power supply, or

$$\frac{C\left(V\Delta V - \frac{\Delta V^2}{2}\right)}{CV\Delta V} = 1 - \frac{1}{2}\left(\frac{\Delta V}{V}\right).$$

When $\Delta V/V = 1$, full discharge efficiency is $1/2$, which we knew. For small per-unit discharge—say, 0.1, or 10% droop—efficiency is $1 - 1/2 \times 0.1 = 1 - 0.05 = 0.95$, or 95%.

In most high-power applications, however, resistance is not used between power supply and capacitor bank. Recharging current from the power supply is always limited by the equivalent ac series impedance of the power supply, most of which is the leakage reactance of the inductively coupled components, especially the high-voltage rectifier transformer.

10.4 Vacuum tubes as switch tubes

Power grid vacuum tubes have numerous characteristics that define their total performance capabilities as pulse-modulator switching devices. The first characteristic of importance is peak-pulse power output, which is the product of the simultaneous maxima of load voltage and current that can be switched. Of course, the maximum load voltage is related to the continuous voltage that can be applied between anode and cathode of the tube while maintaining an accept-




Figure 10-7. How cathode emission density is affected by spacing between grid and cathode (effect on perveance).




Figure 10-8. How current density in region between anode and outermost grid is affected by the spacing between them and effective drive voltage, due to the creation of a virtual cathode between the electrodes.

ably low spark-down rate. Already it is obvious that there will be a degree of flexibility in what the usable anode hold-off voltage really is. In an application where a switch-tube arc is tantamount to disaster, an extremely conservative value must be used that will probably be even less than the manufacturer's suggested operating voltage. Even this number is already somewhat less than the "absolute-maximum independent rating," which defines voltage-hold-off capability at essentially no plate current.


The useful load voltage is affected by load current as well. Maximum load current depends upon the type of cathode, its surface area, and temperature, which is related to heating power; and the current split between anode and other electrodes, such as the grid, shield grid, and screen grid. Although basic cathode research is ongoing and types capable of sustaining high emission densities over long lifetimes have come forth, power-grid-tube cathodes are still of two basic types: thoriated tungsten (directly heated) or oxide-coated (unipotential, or indirectly heated). (Million-hour emission lifetimes can at least be projected from today's data bases for dispenser-type cathodes.) As current from the cathode increases, so must the minimum voltage between anode and outermost grid in order to avoid formation of a "virtual cathode" in the space between the two. Such an occurrence would prevent the passage of electrons through the region. (The electrons will not even leave the cathode unless there is adequate effective drive voltage.)

Figure 10-7 relates cathode emission density to grid-cathode spacing and effective drive voltage. (The last parameter includes the contributions of anode, screen, and control-grid voltages as they are weighted by their respective amplification factors.) Figure 10-8 shows how the current density in the region between the outermost grid and anode is affected by drive voltage and spacing between the anode and outermost grid. These relationships can be defined by the equation

$$j_0 = \frac{2.33 \times 10^{-6} (\sqrt{V_A + V_G})^3}{d^2},$$

where j_0 is the density of current that can traverse the region between the outermost grid and anode, V_A is the voltage between anode and cathode, V_G is the voltage between grid and cathode, and d is the spacing between the outermost grid and anode. Notice that the useful component of current that travels from the grid to the anode is inversely proportional to the square of the distance between them.

Figure 10-9 shows why this is not good news. The figure relates anode-grid hold-off voltage to the spacing between them. The uppermost plotted line, A, is the plane-parallel conductor experimental data, which is of mostly academic interest. Note the intercept at 100 kV for 1-cm spacing. The 100 kV/cm gradient is often cited as a criterion for peak electric field in electron guns that use smooth copper surfaces, such as the ones for large klystrons. (An electric field of 100 kV at 1 cm may sound like a lot, but it is only about three times as great as the corresponding number for the dielectric strength of air at sea level. The actual geometries of power gridded tubes behave in less predictable fashion.) Getting back to Fig. 10-9, element B is a region rather than a line. It defines what one



*Experimental data by Kilpatrick, RSI, Oct. 1957

Figure 10-9. Voltage hold-off between electrodes in vacuum as functions of spacing and surface quality.




Figure 10-10. Typical directly heated thoria-coated tungsten wire-type cathode for power-grid tube.

might expect for a tube with a wire grid and thoria-coated tungsten (ThW) cathode. Region C defines the performance to be expected if the cathode is oxide-coated rather than thoria-coated. Oxide cathodes are, quite literally, flakier than ones made of thoria-coated tungsten, and the products of surface deterioration can contaminate the grid wire, producing localized electric field enhancement. Tubes with higher μ factors have more grid wires that are closer to the cathode so that the wire mesh functions more like an equipotential surface rather than individual wires with respect to the electric-field-strength enhancement produced by conductors of small radius-of-curvature. Smoothness of the wire is also important, but once an arc of sufficient action has terminated on a grid (or screen) wire, that wire, even if it is still there, will no longer be smooth.

If oxide-coated cathodes perform less well with respect to high-voltage hold-off, why are they used at all? The reason is that thoria-coated tungsten cathodes are not as efficient emitters with respect to heater power. For example, on a thoria-coated tungsten cathode, one ampere of peak emission requires more than 15 W of heater power—and this is if a highly regulated filament voltage is used to preclude temperature-limited operation. (Pure tungsten filaments, on the other hand, are often deliberately operated in a temperature-limited fashion in current-control diodes, where current is determined by filament power rather than anode-cathode voltage). There is no pulse-duration limitation for such cathodes.

A comparable oxide-coated cathode may require less than 1/3 as much heater power for the same emission current, but pulse durations are limited to tens of




Figure 10-11. Typical helical wire-type control grid for high-power triode vacuum tube.

microseconds and duty factors to 1% or less, which is more than adequate for many applications. Useful emission density for both types is 2 to 3 A/cm². Even though emission from an oxide cathode can be 10 to 50 times this great for microsecond-length pulses, the only current that matters is what can get from outermost grid to anode. This current is limited by the virtual-cathode effect mentioned earlier.

The thoriated-tungsten filament can be expected to last more than 10,000 hours (or one year of continuous operation). This longevity is comparable to the life expectancy of the rest of the equipment if utilization is sporadic. The life expectancy of oxide-coated cathodes depends on too many factors to be accurately predicted or even repeated, but it is typically less than half that of thoriated-tungsten.

Figure 10-10 shows how a directly heated thoriated-tungsten wire cathode for a high-power gridded tube is arranged. Such structures can become quite tall, and, when excited with ac heater power, can vibrate like tuning forks due to magnetic force caused by the filament current. Such filaments may require dc or high-frequency ac in order to produce no unbalanced magnetic force. (Or, as we will see later, such filaments may actually be built in three symmetrical sections in order to be excited with balanced 3-phase ac.)

Figure 10-11 shows what a similar tube looks like after a helical grid has been wound around the cathode wires. The closer the grid is to the cathode, the

higher the μ factor will be, the higher the cathode-current-density change will be for a given grid-voltage change, and the greater the anode-voltage hold-off is likely to be. The wire-to-wire spacing of the grid, however, will have to be smaller, approximately equal to the spacing between grid and cathode. This is done so the "screening" effect of the grid, or the amount of total grid area that is physically blocked by the presence of grid wire, will be greater than that of a lower- μ grid spaced farther away from the cathode. The greater the screening effect, the more likely the grid is to intercept electrons on their way from cathode to anode when the grid is driven positive with respect to the cathode. This screening effect is reduced if the diameter of the grid wire can be reduced. But the diameter must still be large enough to dissipate heat, including the heating caused by its proximity to the hot cathode and the interception of electrons. Even more importantly, the wire must have cross-sectional area (or AWG gauge) sufficient to absorb the action content of an arc discharge without fusing. (As will be discussed later, the transmitter designer is responsible for limiting the arc action by employing surge-limiting resistance and, when needed, an electronic crowbar charge diverter.) Even in high- μ tubes, the grid screening factor is less than 25%, and it typically varies down to 10% for low- μ tubes. Grid current at anode-saturation voltage is typically 1.5 times the screening factor, so that a 25%-screened, high- μ triode may intercept grid current that is 1.5 times 25%, or 38%, of cathode current, compared with 1.5 times 10%, or 15%, for a low- μ triode.

Grids are almost always wound of tungsten or molybdenum wire that has been coated with gold or platinum. This is done so they can operate at temperatures higher than the cathode—even as high as 1400°C (compared to 1000°C for the cathode)—without becoming primary emitters of electrons themselves. When they are driven positive with respect to the cathode they will attract and intercept electrons. The sum total of these electrons constitutes grid current, and the product of this current and the peak-grid voltage is the peak-grid dissipation. This current and power must be supplied by the external grid-drive circuitry and must be dissipated by the grid as heat.

Because of an effect called secondary emission, not all electrons that collide with the grid are counted as grid current, however. Primary electrons strike the grid surface and knock off others, which continue on to the anode. So long as the secondary emission is less than the primary interception, things look good for both the tube and driver. On the other hand, if more secondary electrons are emitted than primaries are collected, grid current is negative and the grid can run away, unless it is externally swamped by shunt resistance. Pulse stretching is the usual symptom of excessive secondary emission; cathode and plate currents continue even after the grid drive pulse has stopped. Pulse stretching in some thoriated-tungsten tubes can be corrected by a process called grid blackening. In this process, filament voltage is increased to twice its normal value (so filament power approximately trebles). Some of the filament carburization boils off and is deposited on the grid, making its surface stickier and less prone to secondary emission. Secondary emission is less likely from grid wires whose surfaces have been made intentionally rough, but anode-voltage hold-off is degraded as a consequence.

Recently, manufactures have made grids and screens of pyrolytic graphite. This material does not emit secondary electrons, but tubes using it do not always

please their users because net grid current is always higher for them than conventional grids. Higher current means that grid drive circuits and screen-grid power supplies have to be rated for more output current. Nevertheless, with pyrolytic-graphite grids and screens there is less likelihood of tube-to-tube variations, or variations within a single tube over its life span.

In oxide-cathode tubes, pulse duration is limited by the cathode. In thoriated-tungsten tubes, pulse duration is usually limited by intrapulse grid heating, which, ignoring heat lost by radiation, is adiabatic, meaning that no heat is dissipated during the pulse itself. Therefore, the temperature rise is proportional to the pulse energy, which is the product of peak power and pulse duration, and inversely proportional to the product of grid-wire mass and specific heat. If power is expressed in watts, pulse length in seconds, and mass in grams, the proportionality constant for temperature rise in degrees Celsius is 0.24. With grid temperature at 1400°C, however, radiation losses are fortunately not insignificant.

The last electrode of a power grid tube is, of course, the anode, which is where we would like the majority of our cathode-emitted electrons to end up. A tube designer wants the electrons to arrive at the anode with a kinetic energy that is as small as possible relative to the anode area and the spacing between anode and outermost grid. This kinetic energy and the rate of electron arrival—which is, of course, the anode current—determine the anode dissipation. Only so much can be done to minimize it, and a conduction-anode drop that is 10% of the anode hold-off voltage is considered a good design, resulting in 90% anode efficiency as a modulator switch. The bulk of anode design, other than geometric requirements, falls in the category of heat transfer, which will be touched upon in a later section.

Needless to say, grid and/or screen current does no one any good, and tube designers have exercised a great deal of their creativity in minimizing it. Now we will look at a few actual tube designs to see what has been accomplished in this regard. (Remember, of course, that almost nothing new has been done in this field for almost 30 years.)

10.4.1 *The 7560 triode switch tube*

Figure 10-12 shows the plate (anode) current characteristics for the 7560 triode modulator switch tube, which once was simultaneously manufactured by three of the largest and most competent power-vacuum-tube sources in the world. Today it may not be made by anyone, because two of the three have ceased operations and the third may no longer be interested. Nevertheless, there is still a host of operating sockets for the tube type, and shops that can rebuild it still flourish. The figure also shows the characteristics of another power triode, the 8547, which is not primarily a modulator switch tube because its nominal anode-voltage rating is 17 kV. Both tubes have water-cooled anodes capable of dissipating 175 kW of power. The 7560 differs from the 8547 only in the increased spacing between its anode and grid, which is required to achieve its 50-kV anode-voltage hold-off rating. Because the anode is farther away from the grid in the 7560, a designer would expect that greater grid-anode voltage drop would be required for a given amount of anode current given its virtual-cathode effect. Indeed, the effect is clearly illustrated in Fig. 10-12a by the slope of the “diode





Figure 10-12. Voltage/current characteristics for two power triodes: type 7560 (left) and type 8547 (right), which differ only in that the 7560 anode is farther from its grid.

line," which is the line (not exactly straight) that slopes upward and to the right from the origin. It is also the line where all of the anode current lines merge. This diode line quantifies the virtual cathode and has a much steeper slope in the closely spaced electrodes of the 8547 than in the widely spaced ones of the 7560.

If we select operating points for both tubes that meet the arbitrary but sensible 90%-anode-efficiency criterion—at which the average output power could conceivably be as high as 175 kW/0.1, or 1.75 MW—we can determine how well grid current is managed in the two tubes. For the 7560, the conduction drop selected is 5 kV, which is 10% of 50 kV, and for the 8547 it is 1.7 kV, 10% of 17 kV. With positive grid voltage of 2000 V, the anode current of the 7560 is 380 A, and its grid current is 130 A. Therefore, the ratio of the two values is 2.10. Of the total 513-A cathode current, which is the sum of anode and grid currents, 74% became anode current and 26% became grid current. This breakdown suggests a grid-screening fraction of 17%, which would put the 7560 in the medium- μ category. (Its μ , in fact, is 45.) With 50 kV between anode and cathode, the grid must be negative with respect to cathode by approximately 1.1 kV for "projected" plate-current cutoff. (The actual plate-current cutoff requires an additional safety factor of negative grid bias.) The 8547, with a positive grid voltage of 2000 V, an anode current of 390 A, and a grid current of 100 A, has a ratio of 3.10. Of the 490-A cathode current, 80% made it to the anode, and 20% to the grid, suggesting a 13% grid-screening factor and lower μ . We know, however, that the tubes are identical except for anode-grid spacing, so the grid-screening factors are the same. The μ is lower, however, because the anode is closer to the cathode ($\mu = 14$).

A: Average-Plate Characteristics



B: Average Grid Characteristics




Figure 10-13. Voltage/current characteristics for the type 8461 triode.

The filaments of both tubes operate at 14.5 V and 450 A, or 6250 W. At the rate of 15 W per cathode peak ampere, the available cathode current should be 435 A. The actual peak-cathode current rating is 550 A. Although the 8547 more efficiently converts cathode current to useful anode current, its peak-pulse output power under the conditions evaluated is 6 MW, whereas the peak power from the 7560 is almost 20 MW. Neither of these tubes employs any special strategy, either geometric or electromagnetic, to minimize grid current. The grid is no more sophisticated than a helix of wire wound around the array of cathode wires.

10.4.2. *The WL-8461 triode switch tube*

This is another example of a thoriated-tungsten cathode. The WL-8461 is a "big but simple" power-grid triode modulator switch tube. Many sockets throughout the world are also filled with this tube, but it is no longer being manufactured. It is, however, being successfully rebuilt.

The WL-8461 is "big" in that its anode-voltage hold-off rating is 75 kV, its peak cathode-current rating is 750 A, and its water-cooled anode-dissipation rating is 200 kW. It is "simple" in that nothing special was designed into this tube to minimize grid current. Figure 10-13 shows its anode- and grid-current characteristic curves. Using the same 90%-anode-efficiency criterion (which yields a conduction-anode drop of 7.5 kV) and a positive grid voltage of 3.5 kV, the anode and grid currents are 600 A and 92 A, respectively, for a total cathode current of 692 A. The current ratio, therefore, is 6.5. Anode current is 88% of cathode current and grid current 12% of it. These values suggest a low μ . (Its μ is actually 25, which is somewhat low.) In Fig. 10-13a, note the cross-hatched area in the anode (plate) current characteristics just to the right of the virtual-cathode diode line. This cross-hatch area is bordered approximately by the +1500-V and +3000-V grid-voltage lines and the 3-kV and 7-kV anode-voltage lines. This is a region where a designer can expect oscillatory instability when the frequency is in the VHF range. The problem derives from the virtual cathode; there is a region of instability between the grid and anode where electrons don't know whether to go forward or backward. (This same effect was deliberately exploited in the Vircator, or virtual-cathode oscillator, an ultra-high-power microwave device.)

Another dimension of 8461 bigness is its directly heated cathode, which is physically so large that it is divided into three sections and excited with 3-phase ac. Each section operates at 15 V and 200 A, for a total of 9 kW input. At 15 W per ampere, the cathode current rating would be 600 A. Figure 10-14 shows the external physical dimensions of the 8461 and the 7560 to give an idea of the large size of tubes in this class.

10.4.3 *The ML- 6544 triode switch tube*

The objectives in placing the control grid close to the cathode are three-fold:

- to maximize the perveance, or the current density, from the cathode as it relates to the 3/2-power of positive voltage applied between grid and cathode;
- to maximize the μ -factor, which minimizes the negative voltage that must be applied between grid and cathode (or screen-grid and cathode in a




Figure 10-14. Physical attributes of type 8461 and 7560 power-grid triodes.

tetrode) for a given voltage between anode and cathode;

- and to minimize the electric-field enhancement at the anode-facing surface of the grid, which improves the anode-grid voltage-hold-off capability.

The price that must be paid for placing the control grid near the cathode is increased grid heating due to primary electron interception and the proximity of the hot cathode itself.

The type 6544 triode was designed especially for modulator switch-tube service, and it is not a simple triode. Not only does it have a control grid but it has a shield grid as well. In addition, it has an oxide-coated unipotential (indirectly heated) cathode of special shape. All of these features, shown in cross-sectional




Figure 10-15. Section view through active region of type 6544 shield-grid triode.

view in Fig. 10-15, are intended to contribute to a high- μ triode ($\mu = 90$), which carries the implication that it has high grid-current interception, but without the high grid current. The cathode is of cylindrical shape, but its coated active emitting surfaces by no means cover the entire available surface. In fact, as shown, the individual emitting surfaces are groovelike concavities spaced around the periphery. Control-grid rods, held in place by ringlike headers at each end, run the length of the emitting grooves, but they are staggered in spacing with respect to grooves so that they are in between them and not in the line of fire of electrons leaving the cathode on the way to the anode. The shield grid is another




Figure 10-16. View of concentric subassemblies that make up the type 6544 shield-grid triode.

group of rods that has been placed on the same radials as the control-grid rods. The shield rods are located just on the anode side of the control grid on a circle of slightly larger diameter. The shield grid is connected to the cathode. Figure 10-16 shows the normally concentric electrode assemblies lying side by side. Instead of a uniformly emitting cathode surface, the 6544 has a number of individual electrostatically focused electron guns cylindrically arrayed. This geometry "hides" the control grid from the cathode as much as possible in a high- μ tube.

To what extent does this strategy succeed? The anode-voltage rating of the 6544 is 20 kV. Using the same 90%-efficiency criterion, at an anode voltage of 2 kV and a peak-grid voltage of 1300 V, the peak-anode current is 70 A and the grid current is 5.5 A, as shown in the current/voltage characteristics of Fig. 10-17. The ratio is almost 13, which is certainly not bad for a tube with μ of 90. Total cathode current is 75.5 A, which is approximately the peak rating of the tube. Of this, 93% is anode current and 7% is grid current. The filament heating power is 360 W, giving an emission loading of 4.8 W/A.

10.4.4 The S94000E and 4CPW1000KB tetrode switch tubes

The ML-6544 is a triode in that only three electrodes are externally accessible. There is, however, a fourth electrode, the shield grid, but it operates at the same voltage as the cathode and is internally connected to it. A true tetrode has an independently operated fourth electrode, which is physically similar to the 6544 shield-grid and is typically operated at a continuously applied (or sometimes pulsed) voltage positive with respect to the cathode, usually in the range be-




Figure 10-17. Voltage/current characteristics of type 6544 shield-grid triode switch tube.

tween 1000 and 2000 V. This electrode shields; or "screens," the grid from the anode, so it is usually referred to as the screen grid. It is also placed very much closer to the control grid than to the anode. For all intents and purposes, therefore, the screen grid is a virtual anode to the control grid, operating at constant voltage with respect to the cathode. If positive control-grid-to-cathode voltage is required for the desired peak-cathode emission (and it often is not), the amount of current intercepted by the grid for any fixed grid-cathode and screen-cathode voltages will be substantially independent of the anode-cathode voltage. Often, because of the electric-field strength at the cathode that results from screen-cathode voltage, full emission can be obtained with zero or even negative grid-cathode voltage, in which case there will be no grid power dissipation at all due to primary electron interception. However, the grid (and the screen as well) will be heated by radiation from the hot, nearby cathode.

The amplification factor in a tetrode is defined with respect to the screen grid rather than the anode, as is the case with the triode. Tetrode amplification factors range typically from 5 to 10. Projected cathode-current cutoff occurs when a negative control-grid voltage, which is affected only slightly by the anode-cathode voltage, is applied that is between one-tenth to one-fifth of the positive




Figure 10-18. Physical aspects of S94000E modulator switch tube.




Figure 10-19. Voltage/current relationships in S94000E tetrode switch tube for screen voltage of 1000 V.

screen-cathode voltage.

Many tetrode types (and triode types as well) are used as high-power hard-tube modulator switches, although they were not designed for such service. These tetrodes range in capacity from the 100-kW anode dissipation of the 4CW100,000 right up to megawatt-dissipation tubes. They were originally intended for use in high-power broadcast transmitters or other high-power RF sources at frequencies through the HF band (but usually not above 50 MHz). In RF amplifier service, a dc anode voltage rating in the 20-kV range is common. This may seem inadequate for most hard-tube modulator service, and it probably is. The maximum instantaneous anode voltage that an RF amplifier experiences is nearly twice the dc component, however. Moreover, a tube that is rated for high-level amplitude modulation at 100% positive modulation will see a "dc" or audio-frequency peak that is also twice the average anode voltage. (Remember that the peak instantaneous power from a 100%-amplitude-modulated power amplifier is four times the carrier power, and even a 50-kW-rated transmitter produces 200-kW peak envelope power.) It can be seen that in modulator-switch service, where maximum voltage hold-off occurs during nominally zero-current situations, the true voltage-hold-off capability may be many times the RF-amplifier dc anode-voltage rating. This is especially true if the tube is specified at procurement for such service. In such cases, it can be specially processed during manufacture to enhance high-voltage tolerance.

Two good examples of tetrodes designed expressly for service as high-voltage modulator switch tubes and intrapulse voltage regulators—and not just specially processed RF amplifier tubes—are the S94000E and the 4CPW1000KB. The




Figure 10-20. Transfer characteristic for S94000E tetrode showing "projected" cutoff.

S94000E has an anode-hold-off voltage rating of 200 kV and a maximum peak-anode current of 125 A. (We will later see that peak-anode current is virtually the same as the peak-cathode current.) The 4CPW1000KB has an anode-voltage rating of 175 kV and a peak-anode-current rating of 75 A. Both tube types evolved from high-power RF amplifier forebears. Like most tubes intended for high-voltage hold-off capability, they differ from their predecessors in the spacing between anode and screen grid. Consequently, the slope of their diode, or virtual-cathode, lines also differ.

The 4CPW1000KB is a radial-beam, external-anode tetrode of more-or-less conventional geometry. It descended from the 4CW250000 and was intended primarily as a megawatt-class RF amplifier in the HF band. Its anode is capable of 1 MW of average dissipation—which is what the "1000 K" in its type designation refers to—giving it extended capability for linear voltage regulation and waveform shaping.

The S94000E, by comparison, is literally inside out, as shown in Fig. 10-18. It is a radial-beam tube, but its cylindrically shaped anode is on the inside of the tube rather than the outside. Around the outer periphery is a circular array of "unitized" electron guns, which are shown conceptually in the figure. Each gun has its own central cathode bar, grid box, and screen box. The grid and screen boxes are solid on three sides. The sides facing the anode are open, except for the very fine wires stretched across them that are the actual control and screen grids.




Figure 10-21. Area of permissible operation for S94000E tetrode switch tube.

This electron-optics technique first emerged in a grid-driven power tetrode for UHF television transmitter service and resulted in a many-fold increase in the power-frequency product obtainable from power gridded tubes. The voltage/current characteristics of the tube for a constant screen voltage of 1000 V are shown in Fig. 10-19. Note that these curves have a different format than the ones previously shown. These are constant-current characteristics rather than constant-grid-voltage characteristics. They illustrate the almost-total screening effect of the anode from the cathode. This is due to the fact that anode current is all but independent of anode voltage. The effective perveance, or transconductance, of the electron optics is such that the full current is obtained when the grid voltage is never positive with respect to the cathode, and the total "grid base" (the grid-voltage swing from cutoff to full current) is only slightly more than 150 V. This condition requires virtually zero driver power. With the grid never positive, there is no grid current either. The grid that competes with the anode for electrons is the screen. Note that even with an anode voltage that is only 5% of the hold-off rating (10 kV compared to 200 kV), the current-split ratio between screen and anode is 120 A/6 A, or 20, which is superior to the conventional triodes previously discussed.

Figure 10-20 shows how anode current varies with grid voltage for a constant screen voltage of 1000 V and constant anode voltage of 12 kV. This figure illustrates the concept of "projected" cutoff, which is the dashed line. It "projects" a grid voltage for anode-current cutoff of about -110 V. For a screen voltage of 1000 V, this relates to a screen-grid amplification factor, or μ_s , of 1000 V/110 V, or 9, which is the nominal screen μ . Figure 10-21 shows the dimensions of the tube's permissible operating region, which is the equivalent to the safe operating area (SOA) for a field-effect transistor. The left-hand, sloping limit line is the diode




Figure 10-22. Translating tetrode voltage/current characteristics for different screen voltages.

line. To the left of this line screen current dominates anode current, and screen dissipation becomes excessive. The top segment is the maximum anode current limit. The curved-line limit is the actual anode dissipation, beyond which the product of anode voltage and current exceeds 2 MW. The right-hand line, which rises from the 30-kV anode voltage, is a limit on electron energy, or velocity, beyond which electrons will penetrate the anode and cause damage.

The voltage/current characteristics of tetrodes, in general, are more cumbersome to illustrate because there are four degrees of freedom compared to only three in the triode. This means that there is a unique set of data interrelating anode and grid voltages with anode, grid, and screen currents for each and every value of screen voltage. Manufacturer's published data will rarely show more than two such screen voltages. Even if no more than one is given, however, it is possible to synthesize the appropriate data needed for any other value of screen voltage. How this is done is shown in Fig. 10-22. The manufacturer's information given in Fig. 10-22a and is based on constant-current characteristics for 1500-V screen voltage. But what is wanted is the same data for a screen voltage of 750 V. To derive this information, a new set of axes along with new grid voltage and anode voltage must be generated. In this case, the ratio of 750 V to 1500 V is $1/2$. The original current values for grid, screen, and anode are all multiplied by the transfer ratio $1/2$ raised to the $3/2$ power. In the example shown, the original 250-A anode-current line becomes in Fig. 10-22b $250 \text{ A} \times (1/2)^{3/2}$, which is 86 A, and the original 400-A line becomes the 141-A line by the same process. Meanwhile the 25-A screen-current line becomes the 8.8-A line and the 3-A grid-current line becomes a 1.06-A line. An entire characteristic can be translated, but plotting only a few points at the full-conduction and cutoff ends usually suffices because we are using these devices primarily as switches.




Figure 10-23. Construction of the type ML-8549 magnetically beamed hard-tube switch.

These synthesized points will be as useful to the designer as the original published data, which are offered with no guarantee that they characterize any given tube of the type. What, then, guarantees the performance for any tube type, triode or tetrode? The guarantee lies within the specific quality-assurance test specifications for a given tube type. These specifications are in one of two formats: EIA or MIL-E-1. These specifications typically establish ranges of voltages and/or currents within which a particular tube must fall in order to be sold under its part number. For instance, the negative grid voltage required for the plate current to be less than some specified value at a given anode voltage (or a combination of anode and screen voltages) will not be expressed as a single value but as a range of voltages. The transmitter design had better be able to accommodate the entire range or the worst case. In other words, design to the tube specification not the tube data sheets.

10.4.5 The ML-8549 magnetically beamed triode

The tubes discussed so far have used only geometry and electrostatics to affect electron trajectories. The ML-8549 was one of the first attempts to use a combination of geometric, electrostatic, and electromagnetic influences on electron trajectories in order to maximize the number of them that make it to the anode. The type 8549 triode is illustrated in Fig. 10-23. Note the cross-section view of its inter-electrode action space. It is a radial-beam device. Its elements are cylindrical and they are coaxial. The inner- and outermost elements, however, are not electrodes at all. They are magnetic pole pieces whose orientation is similar to those used in dynamic loudspeakers, in that they support radial magnetic field lines. The space where the voice coil of a loudspeaker would fit is occupied by the electrodes.

But this is no simple triode. There is a wire-type directly heated, thoriated-tungsten cathode. But instead of being the innermost electrode, it is the centermost. Inside and outside of it are two sets of wire control grids that are circumferentially




Figure 10-24. Voltage/current relationships in ML-8549 triode.

staggered with respect to the cathode wires so that they are not on the same radial lines as the cathode wires. Inside of the innermost grid wires and outside of the outermost grid wires are an inner and an outer anode. There are, in effect, two triodes sharing the same cathode. Electrons leaving the outer face of the cathode travel outward, and those leaving the inner face travel inward. Moreover, even though they travel past grid wires that are positive with respect to the cathode, and therefore attract them, their paths are constrained by the radial magnetic field to avoid the grid wires and continue on toward the anodes.

How well did all of this work? Figure 10-24 shows the voltage/current characteristics of the 8549. The tube is rated for 65 kV voltage hold-off. Again using the 10%-conduction-voltage drop criterion, or 6.5 kV, we find that the same posi-




Figure 10-25. Type ML-8618 magnetically beamed triode switch tube.

tive grid voltage that produces 1100-A peak-plate current before (about +2700 V), now produces only 10-A peak-grid current. The ratio is 110:1. More than 99% of the cathode current makes it to the anodes. However, the negative grid voltage required for plate-current cutoff at 65 kV anode voltage is 4000 V, meaning that the amplification factor is only slightly more than 15. The total grid swing, therefore, is almost 7000 V. Even so, a peak-pulse output power of almost 65 MW is obtained from a drive power of approximately 100 kW.




Figure 10-26. Voltage/current relationships of ML-8618 triode.

The tube was not electrically small, either. It had filament power of over 14 kW (7.5 V at 1880 A), which is about 12 W/A of pulse-cathode emission, and an anode-average-power-dissipation rating of 500 kW. Another indication of the reduced grid heating is the maximum pulse duration rating of 0.1 second, or 100,000 μ s. (Remember that grid heating, not cathode emission, is the pulse-duration limiting factor in a thoriated-tungsten filament switch tube.)

This tube is spoken of in the past tense because the application for which it was primarily intended—a giant, experimental pulse modulator—was by no means an unqualified success, and the tube type itself suffered by association. The magnetically beamed concept, however, did not die with it.

10.4.6 The ML-8618 magnetically beamed triode

This tube, as shown in Fig. 10-25, is a very close relative of the type 8549. The electromagnetic strategy is the same, but the geometry is planar instead of coaxially cylindrical. The effect on performance is quite similar, as shown in the voltage/current relationships of Fig. 10-26. The tube has been rated at as much as 50 kV anode-voltage hold-off—although it rarely operated above 40 kV. Using the 50-

kV rating and a corresponding conduction-anode-voltage drop of 5 kV, we see that a typical 150-A value of peak-pulse anode current is obtained with a grid voltage that is 1300-V positive with respect to the cathode along with a peak grid current of about 1.4 A. (The maximum peak-cathode-current rating is 200 A.) The current ratio is slightly more than 100, and over 99% of the electrons emitted by the cathode make it to the anodes.

The grid interception is low—so low that there is no real reason to be limited by the 10%-anode-voltage criterion. Indeed, at a conduction anode voltage of 3 kV, we see that the same 150-A peak-anode current is obtained with a grid voltage 1500 V positive and a grid current of 2.4 A, which is still a ratio greater than 60.

Typical of a thoriated-tungsten-cathode tube with low grid-power dissipation, the 8618's pulse duration limit is 10,000 μ s at a duty factor of 6%. The anode, when properly water-cooled, will dissipate up to 80 kW. Unlike the larger 8549, hundreds of 8618s were built and used and they are capable of being




Figure 10-27. Cross-section view of L-5012 Injectron beam-switch tube.

rebuilt.

10.4.7 *The L-5012 and L-5097 Injectron™ beam-switch tubes*

We have seen how effective collinear electric and magnetic fields can be in minimizing grid current in the examples of the previous tubes. The L-5012 Injectron beam-switch tube (BST), the smaller of the two tubes of this type, was the first vacuum switch tube to make a complete break from the conventional wire-type grid structure. Instead, it used a completely solid control electrode. How can electrons get through a solid metal electrode? They can't, of course. But by using crossed rather than collinear magnetic and electric fields, they don't have to.

Figure 10-27 shows a cross-sectional view of the L-5012. Like the ML-6544, this switch tube uses an indirectly heated oxide-coated cathode that is cylindrical, which means that electrons leave it radially. The control electrode appears to the cathode as an external cylindrical anode, which it is. For this reason the control electrode is called a modulating anode rather than a control grid. When the voltage on the modulating anode is positive with respect to the cathode, electrons leave the cathode radially and are accelerated toward the modulating anode. Surrounding the modulating anode, however, is a stack of permanent magnets that are magnetized so as to produce a magnetic field along the axis of symmetry of the tube, which is at a right angle to the electric field between the modulating anode and cathode. This configuration is often referred to as a crossed-field, as in crossed-field amplifiers or oscillators, of which the magnetron is the most famous example.

The axially directed magnetic field acts to bend the electron trajectories to follow the magnetic-field lines (as in a magnetron), which is why this tube is sometimes called a magnetron-injection electron gun. Once headed in the right direction, the electrons fall under the accelerating influence of the electric field produced by the voltage on the beam collector, which is yet another anode, but one that is insulated from both the modulating anode and the cathode. The collector serves the role of the anode in a conventional triode or tetrode. The electric field produced by collector voltage is in the same direction as the magnetic field. The electron beam, therefore, is a cylindrical "hollow" beam. The electric fields produced by voltages on the modulating anode and the collector are at right angles to each other: the field produced by the modulating anode is perpendicular to the cathode surface, and that produced by the collector is parallel to the cathode surface. A designer would expect, therefore, that the collector voltage would have almost no effect on cathode current, or that the cut-off amplification factor, or μ , between modulating anode and collector would be infinite. A designer would also expect that above the diode line, the resistance of the incremental collector, or anode, would be nearly infinite as well. (This resistance is the amount of collector-voltage change for a given change in collector current, other things remaining constant.)

Figure 10-28 shows the voltage/current relationships of the L-5012 beam-switch tube. If one ignores the voltage scales, the shape of the characteristics is reminiscent of nothing more than the characteristics of a field-effect transistor. Indeed, the L-5012 is a true field-effect device. With no modulating-anode voltage, there




Figure 10-28. Voltage/current relationships of L-5012 beam-switch tube.

is virtually no collector current, regardless of collector voltage. This implies that μ is all but infinite. There is, however, competition for cathode-emitted electrons between the collector and modulating anode when their voltages approach one another, an example of the virtual-cathode effect. If we apply the 10% voltage criterion, we see that for collector current of 22 A and collector voltage of 15 kV (10% of the 150 kV hold-off rating), the modulating-anode voltage is 8 kV and its current is 0.1 A. The current ratio is 220—the highest so far. Above this “knee,” the plotted line of the slope of the collector-current line corresponding to a given modulating-anode voltage is virtually horizontal, denoting nearly infinite collector (or anode) incremental resistance. The practical significance of this will be discussed later.

The L-5097 is very similar to the L-5012 in design and construction, but it is bigger and more powerful. It is rated at 180 kV maximum collector voltage and 70 A collector current, with 60-kW collector dissipation (compared with 150 kV,

30 A, and 12 kW for the L-5012). When operated at its full capability, or in parallel with other tubes in close physical proximity, the use of a solenoid electromagnet rather than a permanent-magnet stack has been found to give superior performance.

The type of construction used in both tube types has another advantage over conventional wire-grid construction that can be crucial. It has to do with what is likely to happen in the event of an internal arc. The cathode surface, for one thing, is not in the line-of-fire of a collector arc (or of back-streaming positive ions, either). The concave circular button on top of the cathode structure is the target for both. There are no fine wires to be fused by the i^2dt integral of the arc discharge. It is this susceptibility in conventional grid-control tubes that engendered the design of the electronic crowbar in the first place. (Crowbars themselves will be discussed in some detail later.) Ultimately, as the current-handling capability of a particular tube increases, it will be the amount of arc current (or the discharge action, $\int i^2dt$) that it can tolerate that will determine its ultimate internal impedance. We would like to be as close to zero as possible, as has been mentioned many times. Any linear resistance that must be added in order to limit arc or fault current is directly in series with the resistance represented by the tube conduction-voltage drop divided by the output current.

It would not be an unfair statement to say that the Injectron type of beam-switch tube represents the state-of-the-art in hard-tube modulator switches. It is sad to note that it is by no means a recent development. The first such tubes were used before 1960 (and the breakthrough that gave us the unit electron guns of the S94000E switch tube took place in the 1950s). There are, however, more recent design concepts that make use of the far more advanced technology of high-power microwave tubes. These devices promise the same type of performance as the Injectron without the use of magnetic fields. The devices, however, have yet to be built.

10.4.8 The 8454H Crossatron™ switch

Although this device is by no means a hard tube, it is not a thyatron either—although it is not coincidental that their names share the last five letters. The first five letters refer to the crossing of magnetic and electric fields, as with the Injectron. This device could be characterized as the “missing link” between the hard tube and the soft tube. Like the thyatron, it is filled with low-pressure hydrogen (0.05 torr) so it is a long way from being “hard” (10^{-8} torr). But it is not truly “soft” like a thyatron, whose hydrogen pressure is usually 0.5 torr. Its conduction characteristics are also intermediate. Whereas a hard tube may have conduction-voltage drop of approximately 10% of voltage hold-off (or 3 kV to 15 kV) and a thyatron may have 1% voltage drop (or 0.2 kV to 2 kV), the Crossatron voltage drop is more likely to be 2.5% to 3%.

The hard tube has proportional and reciprocal full control of output current (proportionality is retained whether current is being increased or decreased). The thyatron has half-control only. (Actually less than half-control, if you consider the fact that, once turned on, a thyatron has no ability to control current amplitude.) The 8454H Crossatron, on the other hand, can be gated into conduction and then gated off. A positive pulse applied to its control grid will turn it




Figure 10-29. Cross-section view of Crossatron on/off control switch.

on. A negative pulse applied to the same grid will turn it off, much in the same fashion as a gate-turn-off (GTO) thyristor. The amount of current that can be gated off is unlikely to be as much as that which can be gated on, unless the grid construction of the tube is deliberately altered to optimize turn-off instead of turn-on.

The ability of the device to emulate full-control—or at least 2/3 control, if one considers intrapulse-current-amplitude control as being worth 1/3—derives from the action of the crossed magnetic and electric fields. Figure 10-29 shows the cross section of a typical Crossatron. Its design is coaxial-cylindrical. It has an internal anode, its current flow is radial, and it has a cold molybdenum cathode and two grids. A continuous voltage is applied to the “source” grid, which stimulates a continuous “source” plasma. The magnets, which surround the tube, are radially magnetized. They produce alternating north and south poles that progress vertically up the tube and a magnetic field that matches the shape of the source plasma. A positive-going trigger pulse applied to the control grid attracts plasma into the region between control grid and anode and initiates full conduction between cathode and anode. A negative-going pulse applied to the control grid will reverse the process.




Figure 10-30. Load-line representation of modulator/load interaction.

If one of these devices is used to supplant a vacuum tube as the switch in a hard-tube modulator, external resistance must be provided to limit current in case of a shoot-through event. As with a gridded vacuum tube, the Crossatron is almost always more susceptible to damage than the diode-type microwave tube it is pulse modulating. For this reason, it determines the magnitude of the series resistance. The value of this resistance, in turn, often determines the usable output power from the switch because its resistance value can easily be greater than the ratio of conduction-tube drop and output current.

Crossatron tubes have been built with as much as 100 kV hold-off and closing currents as great as 12,000 A. When optimized as either a closing-only or opening-only switch, the 8454H is rated at 50-kV hold-off, 2500-A closing, and 1000-A opening. The Crossatron is one of the few full-control (or at least self-commutation) switches that is still in the process of being developed.

10.5 Load lines

Although we have discussed and illustrated the voltage/current characteristics of a number of triode and tetrode modulator switch tubes, we have yet to relate them to the complementary aspects of the non-linear load impedances that they will be expected to drive. Figure 10-30 shows how two types of modulators would interact with three types of loads, described as load lines. In this representation, load voltage/current relationships are plotted starting with the 0/0 points of the voltage/current axes. The modulator, or source, voltage/current relationships are plotted backwards, starting at the open-circuit-voltage, zero-current points.

Of the three load characteristics, only one represents an actual device. (It is also the least linear one.) It is a backward-wave, crossed-field amplifier, the QK-622 Amplitron. As voltage is increased across the QK-622, current builds up slowly, reaching only 4 A at 40 kV, which is near the knee of its characteristic. Between 40 kV and its operating voltage of 50 kV, current dramatically increases, from 4 A to 32 A. At its operating point, its equivalent linear-resistance impedance, its second characteristic, would be 1563 ohms, which is the slope of a straight line from the operating point back to the origin. However, its dynamic, or incremental, impedance is the slope of the characteristic, or $\Delta V/\Delta I$, as it passes through the operating point, which is 200 ohms, not 1563 ohms. The third load characteristic illustrates is a hypothetical linear-beam device whose operating point is the same as the QK-622 Amplitron, and it follows the 3/2-power voltage-current relationship. Its incremental impedance at the operating point would be $\Delta V/\Delta I$, or 1024 ohms.

Also shown are the source characteristics for two types of pulse modulators: line-type and hard-tube. The line-type modulator characteristic starts at 100 kV and zero current at the right side of the graph and ends at zero voltage and 64 A at the left side, passing through its matched-load operating point of 50 kV and 32 A. The open-circuit voltage is twice the matched-load voltage and the short-circuit current is twice the matched-load current, which is a sure-fire criterion for recognizing a matched source. (It is *not*, however, matched to the incremental impedance of either of our microwave-tube loads.) The hard-tube modulator characteristic starts at 60 kV and zero current at the right side. It is discontinuous, however, having a knee. In fact, it is much like the Amplitron characteristic only rotated by 90°. This means that the hard-tube modulator is the Amplitron's dual. As the voltage drop across the hard-tube switch increases beyond the knee, which occurs at 50 kV (10-kV tube drop, working from right to left), the current tends to be essentially constant, which is the exact opposite of what happened to the Amplitron. Therefore, the hard-tube switch has two distinct incremental impedances: one relatively low "below the knee" and another relatively high "above the knee." The most dramatic difference between the two was exhibited by the




Figure 10-31. How modulator internal dynamic impedance interacts with load impedance to attenuate input voltage variations.

Injectron beam-switch tube. The incremental above-the-knee impedance of the L-5097, for instance, is approximately 75,000 ohms, which is what has been illustrated. The below-the-knee incremental impedance is 10,000 V/32 A, or a little more than 300 ohms. The ratio of the two is almost 250:1.

One of the functions of the load-line/source-line representation is to visualize what happens when one or another parameter varies. If, for instance, the charac-

teristic impedance of the line-type modulator PFN were changed, the slope of its characteristic would change as would the point of intersection with a load characteristic, thus establishing a new operating point. Lower impedance would mean higher short-circuit current, tilting the line upward from its 100-kV starting point. Higher impedance would tilt it downward. What is more likely to happen, however, is a change in the supply voltage, which can take the form of imperfect regulation, intrapulse droop, or power-supply ripple. The effect of supply-voltage change, or modulation, is to shift the starting points of the source characteristics either to the right or to the left, moving the lines with them. For ripplelike modulation of source voltage that has no average value, the average operating point does not change either.

The change in load voltage will always be some fraction of the change in source voltage because of the voltage-divider action of the incremental source and load impedances in series. For example, take the case of a line-type source matched to the linear resistor. In this case, the voltage-attenuation factor is two because source and load resistances are equal. Less obvious is the case with the Amplitron load, which has an incremental resistance (200 ohms) that differs from its zero-intercept impedance (1563 ohms). The attenuation factor is the sum of 1563 ohms plus 200 ohms divided by 200 ohms, or about 10. A 1-V change in source voltage produces only a 0.11-V change in load voltage. With the Amplitron driven by a hard-tube source, an increase in source voltage of 1 V—which is in the above-the-knee direction for the hard tube—is attenuated by a factor of $(75,000 \text{ ohms} + 200 \text{ ohms})/200 \text{ ohms}$, or 376. So the Amplitron voltage increases by only 0.003 V. (This is why constant-current hard-tube modulators are used with crossed-field amplifiers for maximum performance.) If, however, the source voltage had been reduced by 1 V, the Amplitron voltage change would have been attenuated by the below-the-knee source resistance of 300 ohms, giving a factor of $(300 \text{ ohms} + 200 \text{ ohms})/200 \text{ ohms}$, or 2.5, which is a long way from 375. Needless to say, the actual knee of the hard-tube characteristic is never this abrupt, but the trend is the same. Figure 10-31 shows the transfer characteristics for a number of source and load combinations.

10.6 Pulse fall time of hard-tube modulator

The current available from a hard tube used as a series modulator switch (and illustrated again in Fig. 10-32) determines the rise time of the voltage pulse applied to the parallel combination of diode-electron-gun load and stray capacitance, which includes that of the switch tube itself. If the switch tube is a power triode, there will likely be greater anode current available when there is full supply voltage across it than when it has reached its intrapulse operating point (typically a 10% tube drop). The higher the amplification factor of the tube, the smaller the difference in currents is likely to be, and in the case of the Injectron-type beam-switch tube, there is virtually no difference at all. The switch tube in this case is a constant-current source, regardless of the voltage across it. Some of its current is required to charge up the stray, or distributed, load capacitance, but at any state of charge, the remainder of the current, which is proportional to the $3/2$ power of the voltage that we are attempting to build up, will flow into the diode-electron-gun load. Therefore, there is no simple circuit model such as an

RC time-constant, because R varies continuously with voltage.

The rise-and-fall waveforms for a constant-current source are identical but inverted, as shown in Fig. 10-32. They have variable slopes like RC time-constant exponentials, except that the closer the exponentials come to the asymptotes, the greater are the instantaneous values of R and the more slowly they approach asymptotes, much like exaggerated RC time-constants. It is possible, however, to evaluate the length of time required for the pulse voltage to change from one normalized value to another by using the formula shown in the equation

$$T_F = t_2 - t_1 = \frac{2C_s}{k\sqrt{V_k}} \left(\frac{1}{\sqrt{a_2}} - \frac{1}{\sqrt{a_1}} \right),$$

where a_1 and a_2 are two such normalized values, typically 0.1 and 0.9. The three important parameters are the perveance of the diode gun, k ; the flat-top pulse cathode voltage, V_k , and the value of the stray capacitance, C_s . In the example shown, a tube that operates at 50-kV beam voltage and 1-A beam current and has a perveance of 0.9×10^{-6} , has a stray capacitance of 200 pF. Given these parameters, the time between the 10% and 90% points of the rise-and-fall waveforms would be 38 μ s if there was a constant-current source of 1 A. This would be unacceptably long for many, if not all, applications.

However, this is a rare case because most cathode-pulsed tubes have higher perveances and operating voltages. Moreover, switch tubes are usually overdriven during the rise time to provide an excess of current in order to charge the stray

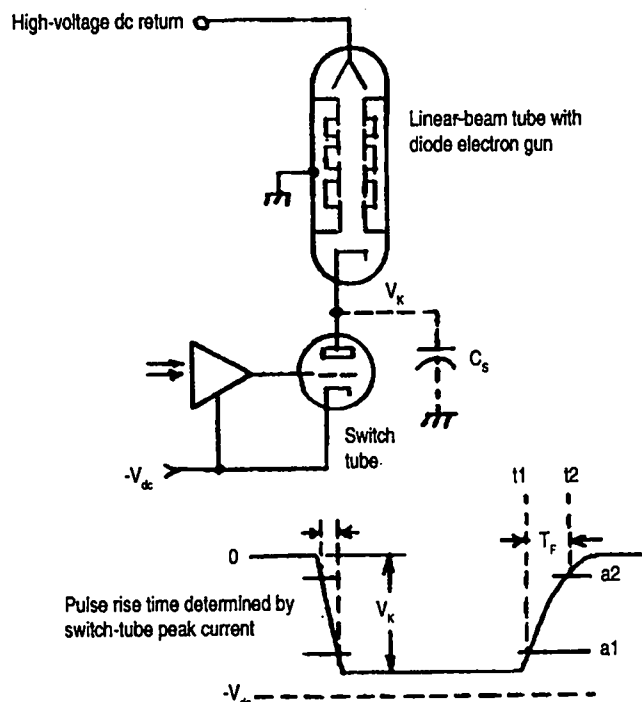


Figure 10-32. Determination of pulse fall time of a hard-tube modulator.

capacitance. The fall time, however, cannot be reduced by anything that can be done to or with the switch tube. When shorter fall time is required, a "tail-biter" or "active pull-up" stage is usually employed. (This feature was illustrated in Fig. 2-1, the transmitter block diagram.) This device, usually another switch tube, is connected between the diode-gun cathode and ground and is pulsed into conduction only during the desired fall time of the pulse. Its current adds to the ever-diminishing current of the diode load to facilitate discharging of the stray capacitance. A constant-current tail-biter of current I will produce a fall time, Δt , that is no longer than $\Delta V \times C_s / I$, where ΔV is the pulse voltage amplitude, even if the diode-load current is negligibly small.

In the case of grid-driven triode switch tubes, which include the modulating-anode type of tube as well, there is another aspect to rise-and-fall time: the effect of "Miller" capacitance, which is related to the capacitance between anode and control grid. Miller capacitance and the capacitance between grid and cathode make up the static input capacitance of a triode. With the anode connected to a fixed voltage source, the rate at which the voltage between grid and cathode can be varied depends upon the driver stage's source resistance, or current-delivery capability, and the sum of the two input capacitance components. The anode, however, is usually not connected to a fixed voltage. It is connected to the load, and the load-voltage excursion is the same as the anode-voltage excursion. The change in charge, or ΔQ , experienced by the anode-grid capacitance is not the product of the grid-voltage change times anode-grid capacitance but is the anode-voltage change times the anode-grid capacitance. In effect, the anode-grid capacitance component of total triode input capacitance is multiplied by the voltage gain of the triode. This fact must be considered in the design of a grid driver. Moreover, the internal impedance of the driver stage in both directions must be considered. This is because all amplifiers that have a single active device have source impedances that are different, depending on whether the device is a source or a sink for output current. For this reason, active pull-up/active pull-down pairs connected in a "totem-pole" configuration are the preferred grid driver stages in critical rise-and-fall-time applications. This is because they can be designed to deliver equal output currents in both directions.

With high-voltage and high-current transistors now available—especially the metal-oxide semiconductor field-effect transistor (MOSFET)—triodes do not have to be grid-driven. They can be cathode driven by what is often called the grounded-grid, or grid-separation, connection. The control grid can be connected to a fixed source of positive grid bias, if necessary. A typical cathode-drive connection is shown in Fig. 10-33. During the switch tube's interpulse (or off) state, the MOSFET is in the off state. To reduce cathode current to the same value as the MOSFET's drain-leakage current, which is negligibly small, the drain of the MOSFET will automatically assume whatever positive voltage is required, including an amount equal to whatever positive voltage is applied to the grid. No negative grid bias supply is needed at all. To cause output current and load voltage, the cathode MOSFET is turned on, pulling the switch-tube cathode to within a few volts of ground (or deck) potential. The effective grid-cathode voltage at this time is whatever positive voltage source is connected between grid and ground. The "Miller" charge and discharge currents flow either into the

ground-reference (or deck-reference) conductor or into the output capacitance of the positive grid-voltage source. This grid-voltage source can be made sufficiently large so that ΔQ through the anode-grid capacitance produces negligible power-supply output-capacitance voltage change.

"Miller" capacitance affects a tetrode-type switch tube about the same way it affects a cathode-driven triode: the electrode closest to the anode is also connected to a fixed source of voltage that is the high-frequency equivalent of ground (deck). The control-grid, or input, capacitance comprises both grid-cathode and grid-screen components, but the screen appears to the grid as an anode connected to a fixed voltage source. The screen-voltage supply, however, must act as both a sink and source for the ΔQ produced by the anode-screen capacitance charge and discharge without generating significant terminal-voltage change.




Figure 10-33. Cathode drive or "grounded-grid" operation of triode modulator switch tube.

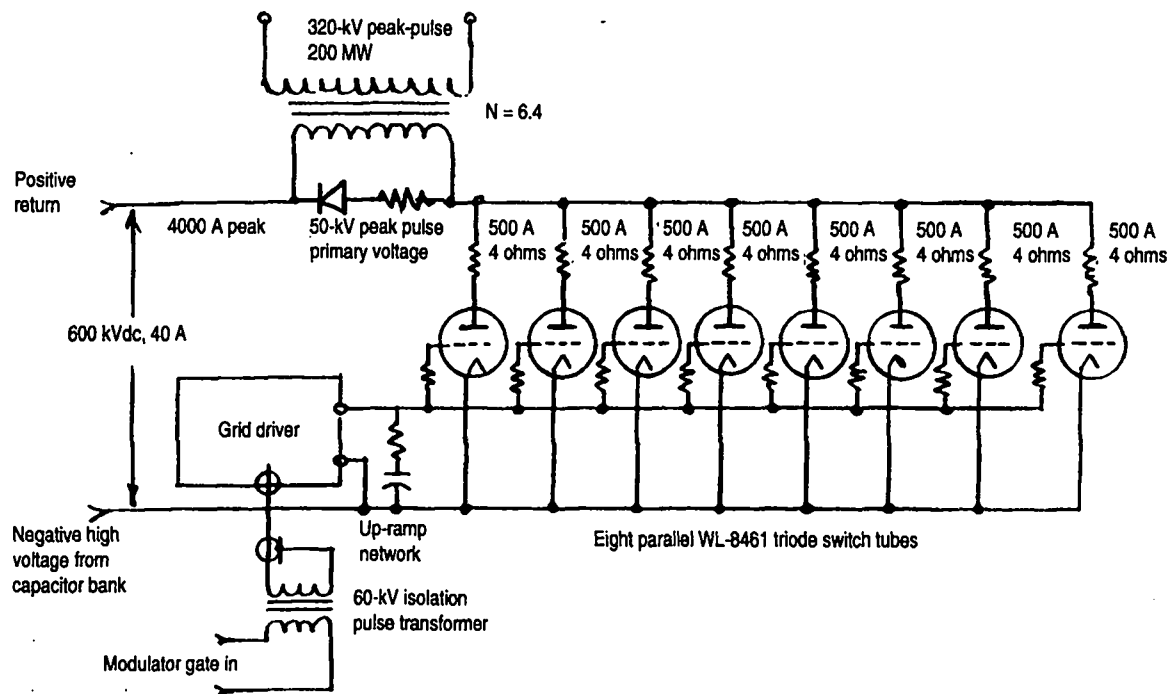


Figure 10-34. Simplified schematic diagram of "200-MW" modulator using WL-8461 triodes.

10.7 Some practical examples of hard-tube modulators

Now let's look at some applications of a couple of these switches.

10.7.1 The 200-MW triode (WL-8461) hard-tube modulator

A hard-tube pulse modulator was actually built using eight parallel-connected WL-8461 triode switch tubes coupled to a super-power klystron load through a 320-kV output-pulse transformer. Its simplified schematic diagram is shown in Fig. 10-34. Each of the triodes was to be operated with peak plate current of 500 A (for a total primary-loop pulse current of 4000 A) and a supply voltage of 65 kVdc, which is adequate to support a 50-kV primary voltage output pulse with as much as 15-kV switch-tube anode-voltage drop.

Although several identifiable pulses at the 200-MW peak-power level were actually obtained, the modulator was never capable of continuous, repetitive operation at this level. The operation of the electronic crowbar in response to individual switch-tube arcs was far too frequent an occurrence, even at output levels below 100-MW peak. This was an experimental modulator, however, and much was learned from it. One empirical lesson, which has not added much to an analytically satisfying data base, was that the performance demonstrated by a single tube operating alone could not be repeated when the tube was operated in parallel with others trying to do the same thing. It might be expected, for instance, that if an individual tube had an average period between internal arc events called T , an ensemble of n such tubes in parallel might have an average spark-down interval of T/n . The observed interval was much shorter than that.

Internal switch-tube arcs are presumably random, non-synchronous events. Extensive diagnostic checks were used, including broad-band delay lines to de-




Figure 10-35. Physical arrangement of components of 200-MW hard-tube modulator (end view: parallel-connected components extend into the page).

lay the pulse output information so that the conditions immediately preceding an arc could be captured on a conventional storage oscilloscope that was triggered by the arc event itself. (This is a data-collection problem ideally suited to today's continuous-sampling digital storage oscilloscopes.) Unfortunately, the results showed that the vast majority of arcs occurred precisely at the end of a current pulse, when the switch tubes were being switched to the voltage-blocking state. So common were these that the acronym EOPSTA was used as log-book shorthand for "end-of-pulse switch-tube arc."

The first experimental objective to become a casualty of this problem was pulse fall time, which was originally intended to be $0.1 \mu\text{s}$. Of course, the culprit




Figure 10-36. Experimental use of low-impedance strip transmission line to reduce inductance of primary-current loop of 200-MW hard-tube modulator.

in such a low-impedance primary loop was series inductance. The inductance of the as-built primary loop—which included the capacitor bank, the eight switch tubes, the electronic crowbar connection, and the pulse-transformer primary connections (see Fig. 10-35)—was $1.7 \mu\text{H}$. When the primary-loop current of 4000 A is brought to zero in $0.1 \mu\text{s}$, or 10^{-7} seconds, the rate-of-change of current is 4×10^{10} A/s, and the impulse, or $L di/dt$, voltage across $1.7 \mu\text{H}$ would be 68 kV, which more than doubled the instantaneous anode voltage. This factor, of course, came in addition to the voltage spike produced by the pulse-transformer back-swing clipper circuit, which shunted the primary winding. Substantial dv/dt snubbing was required before this component of stress was brought to more reasonable levels.

Before giving up on fall time, however, engineers made another attempt to further reduce the primary-loop inductance. A low-impedance, parallel-plate transmission line was used. It had a very large width-to-thickness ratio to make the interconnections, as illustrated in Fig. 10-36. The intent was good but the implementation was doomed to failure because of the electric-field enhancement

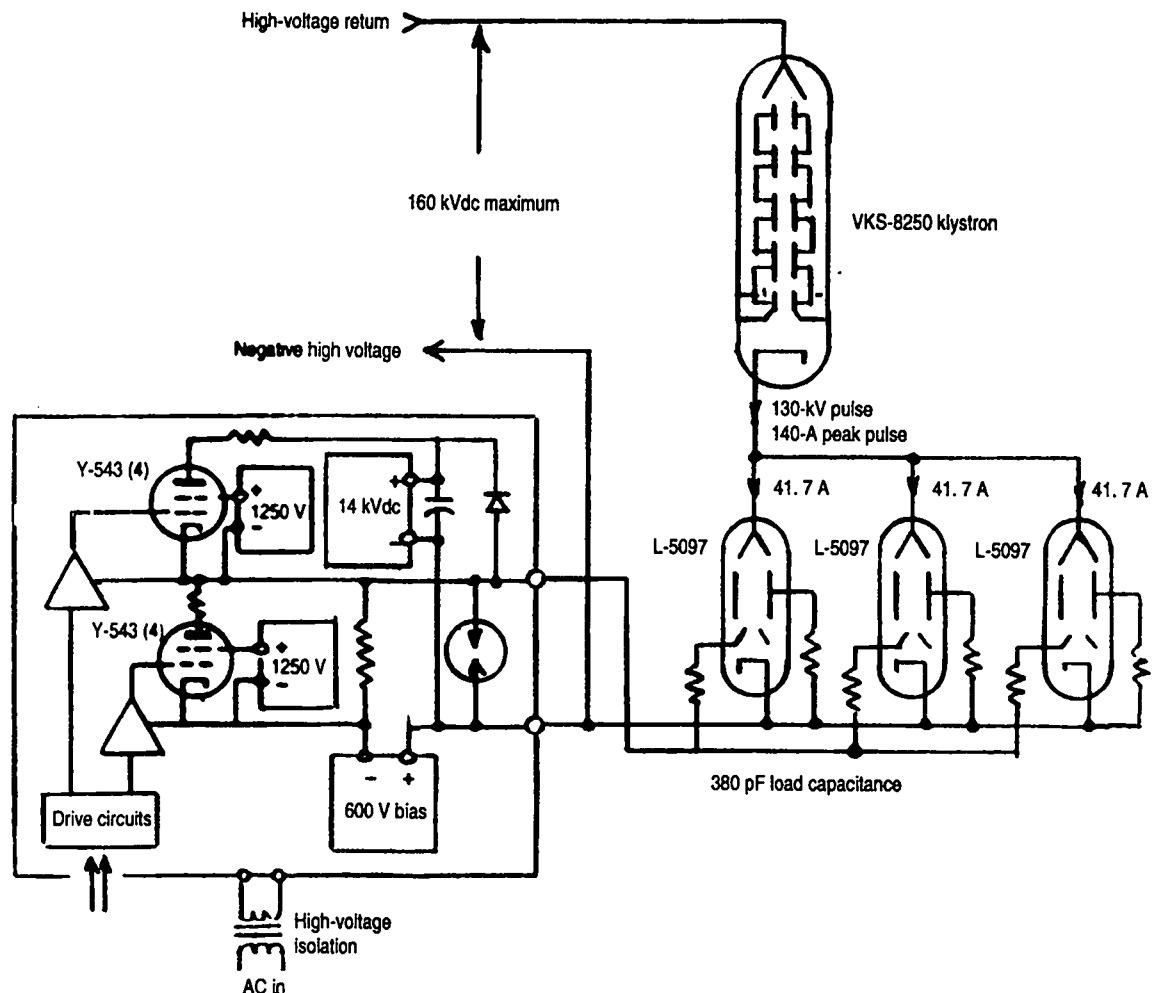


Figure 10-37. Simplified schematic diagram of 18-MW peak-power, high-precision, hard-tube modulator using L-5097 BSTs.

Table 10-1. Operating parameters of 18-MW hard-tube modulator shown in Fig. 10-37.

Parameter	Value
Pulse voltage to klystron	130 kV
Pulse output current	140 A
Peak-power output	18 MW
Pulse duration	3-28 μ s
Burst mode	32 pulses, max., 4- μ s duration, 7-28- μ s spacing
Pulse-pair mode	10- μ s duration, 13- μ s spacing
Rise time	1.2 μ s
Duty factor	0.03, max.
Pulse-top ripple	0.1% (during RF)
Pulse-to-pulse variations	0.02%

at the side-to-side endings of the parallel plates. Merely extending the dielectric material between the plates far beyond the ends of the plates did nothing to prevent corona discharges, which rather quickly punctured the dielectric material.

Reducing fall-time di/dt did not completely cure EOPSTA, however. Other effects were also at work. The temperature of the control grids, which were subject to power dissipation throughout each output pulse, reached transient maxima at the end of each pulse. This made them more likely to emit secondary electrons and less likely to establish zero-current conditions within the tubes when their voltages were brought to the negative voltage that should have been adequate for current cutoff. Experiments were performed where grid drive voltage and conduction-anode-voltage drop were interchanged between so-called hard-drive and soft-drive conditions. Unfortunately, reducing intrapulse grid drive voltage and increasing anode supply voltage to compensate for it produced a statistically insignificant difference in spark-down rate. The only thing that really worked was operating the circuit at lower voltage and lower power.

The suggestion that the problem may have been caused by the tubes' "hot grids" gave rise to an experiment that replaced the eight triodes with eight tetrodes of roughly comparable capability (type 4CW250,000). It turned out, however, that "hot screens" produced much the same effect as "hot grids." Virtually no improvement was noted.

This modulator experiment did, however, lead to a successful 60-MW pulse modulator, which used six parallel-connected WL-8461 triodes that pulsed a 20-MW-peak-power UHF klystron in the transmitter of an instrumentation radar. The system worked well for many years.

10.7.2 An 18-MW hard-tube modulator using the L-5097 beam-switch tube

Waveform versatility and both pulse-top and pulse-to-pulse precision are the salient features of the hard-tube modulator shown in Fig. 10-37. A broad-band high-performance S-band klystron, the VKS-8250, is directly cathode pulsed by a group of three paralleled L-5097 BSTs. The pulsed cathode voltage, which is 130 kV, is driven by a power supply with maximum dc output of 160 kV. The total

klystron beam current of 130 A peak is shared equally among the three L-5097 switch tubes, 41.7 A each. They share current without external isolation components in quite the same fashion as MOSFETs.

The modulating anodes of the three switch tubes are driven from a common pulse source that is configured as a "floating-deck" modulator (which is the next general category of modulator to be considered). This design satisfies a previously mentioned requirement: it behaves as both a source and a sink for the BST's modulating anode displacement currents, which are magnified by the "Miller" effect. Both driver stages use four parallel-connected, high-performance, pulse-rated tetrodes (type Y-543), operating with 1250-V screen supplies. The "pull-up" stage is a "bootstrap" circuit, which means that the common connection for screen and filament supplies and for the low-level drive circuit is the pulse-output bus, which is, so to speak, pulled up by its bootstraps. The driver output is capable of pulsing the modulating anodes of the BSTs to 14-kV positive and returning them to an interpulse negative bias level of -600 V (which is not an absolute necessity for zero BST current).

As shown in Table 10-1 by the listed ensemble of pulse waveforms, including doublets and multi-pulse bursts, advantage is clearly being taken of the electronic agility of the full-control hard-tube modulator switch. Pulse-top variations are 0.1% and pulse-to-pulse variations are 0.02%. Both of these performance achievements are facilitated through the attenuation of power-supply variations by the relatively high incremental-source resistance of the BSTs. The source resistance of each one is 75,000 ohms, giving a parallel-combination resistance of 25,000 ohms for the three tubes. The operating-point incremental resistance of the klystron, $\Delta E_k/\Delta I_k$, is 640 ohms. Therefore, the voltage-attenuation factor is 25,000 ohms divided by 640 ohms, or 40, which is 32 dB.

This is but one of a number of highly successful, highly precise pulse modulators that share very similar circuitry but use different numbers of parallel BSTs.

11. Modulating-Anode Pulsers

The next family of pulse modulators is used with microwave tubes that have isolated, hollow anodes called modulating anodes (see Fig. 8-2). These pulsers must produce output voltages that are substantial percentages of the microwave-tube beam voltage—typically between 50% and 100%—but small intrapulse output current—typically only 0.1% of the total peak-beam current of the microwave tube. Electron guns for modulating anodes requiring pulse voltage in the 50% range are often referred to as “velocity-jump” guns because, even though the voltage between the modulating anode and the cathode determines the beam’s current density, the voltage between the cathode and beam tunnel (or grounded anode) determines the final electron velocity. The electrons, therefore, experience a “velocity-jump” in their passage between the modulating anode and beam tunnel.

The predominant load impedance on the modulating-anode pulser is the distributed capacitance, part of which derives from the capacitance within the electron gun between the modulating anode and the other electrodes, and the rest is associated with the pulse modulator itself, especially if it is of the “floating-deck” type. Modulators of this type have been built having a total load capacitance as high as 1 nF, although a typical value is less than half of that.

Microwave tubes with modulating anodes are designed to operate with voltage applied continuously to the cathode. Such tubes have operated with as much as 150 kV of cathode voltage, although recent technology has limited this figure, more or less, to 100 kV for reliable operation. This limitation, of course, still requires a pulse-modulator output-voltage capability of comparable magnitude.

11.1 Modulating-anode-pulsers topologies

11.1.1 *Passive pull-up, active pull-down*

The simplest of all modulator topologies is that shown in Fig. 11-1. It has a single active switch for pull-down, and it has a resistor, R , for passive pull-up. If the load (the electron gun and modulating anode) requires a full-voltage pulse, it is called a unity- μ modulating anode. In this case, the top of resistor R is simply connected to the high-voltage return bus. If adjustable output voltage over a range of operating pulse widths and repetition rates is required, an adjustable power supply is often used, as shown. This connection qualifies as the simplest because the deck containing the low-level electronics and the common for the switch-tube filament supply are both referenced to a nominally fixed voltage: the negative high voltage required by the microwave-tube electron-gun cathode. The common-mode rejection problem for the low-level drive signal is not as great as it will become in other topologies, as we shall see later.

This circuit arrangement has another modest advantage. The load capacitance contributed by the modulator itself is only the output capacitance of the switch tube, not the entire deck. The disadvantage, however, is usually a crush-

ing one: the off-deck switch tube is in the conducting state during the interpulse interval, which means that there is current flow in the pull-up resistor, R . The value of R determines the pulse rise time, which has the time-constant RC . For instance, if the load capacitance were 500 pF, which is not uncommon, and a time-constant of 10 μ s were desired, R would be 10^{-5} seconds divided by 5×10^{-10} farads, or 20,000 ohms. If the microwave tube operated at 100 kV, the current through the resistor during the interpulse interval would be 100 kV/20,000 ohms, or 5 A. The peak dissipation would be 500 kW. If the duty factor was much less than 10%, the average dissipation would be 500 kW as well. If the pulse duty factor were considerably greater than 50%, this disadvantage would not be a problem. But such a high duty factor is rarely, if ever, the case. Moreover, if the adjustable-output power supply is used, this current must come from it. As rise time is sacrificed, dissipation is reduced, but it is obvious that much needs to be done to make this circuit practical.

11.1.2 Active pull-up, passive pull-down

For the reasons discussed above, the topology shown in Fig. 11-2 is considerably more popular. It also uses a single electronic switch as an active pull-up, with passive pull-down being provided by resistor R . But here, the output pulse is taken from the cathode of the switch tube rather than the anode, as was the case in Fig. 11-1. The switch tube is not a cathode follower, however. Its connection is often called a bootstrap circuit, in that the low-level grid-drive circuits are referenced to the switch-tube cathode, which is connected to an enclosure referred to as a "pulse deck," a "hot deck," or, most commonly, a "floating deck,"

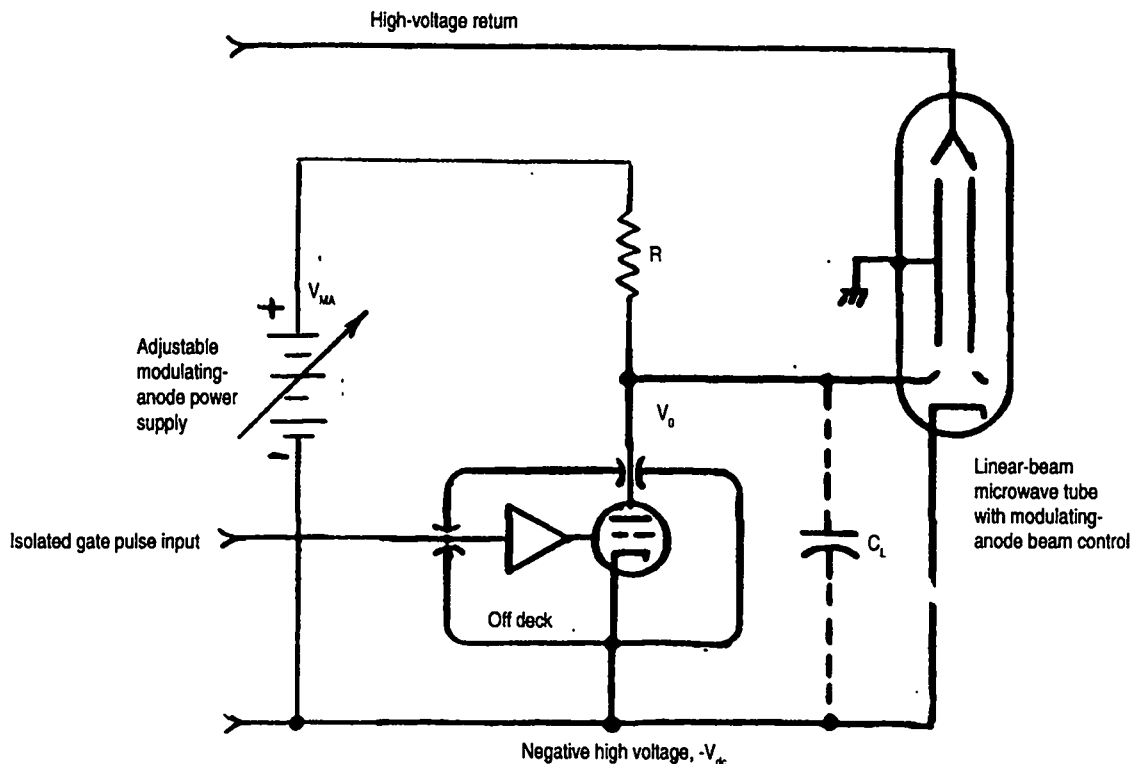


Figure 11-1. Modulating-anode pulser with passive pull-up and active pull-down.

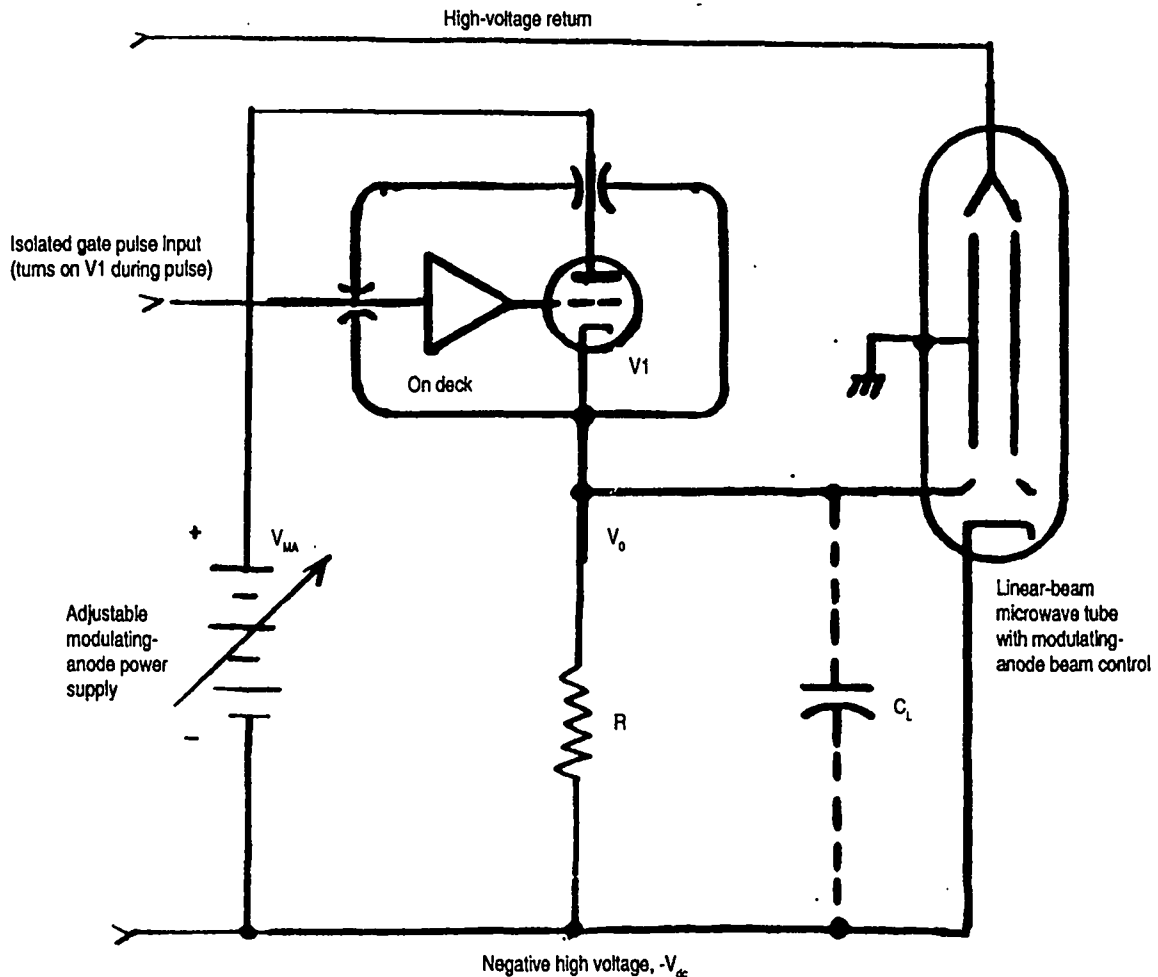


Figure 11-2. Modulating-anode pulser with active pull-up and passive pull-down.

(hence the term "floating-deck" modulator). The high-voltage pulse output is taken from this deck. The low-level circuits, therefore, are pulled up by their bootstraps, so to speak.

The pull-down resistor R now determines the pulse fall time. Its peak current may still be relatively high, as in the previous example, but the average current and dissipation will be proportional to the duty factor, which is seldom greater than 10% and is often less than 1%. However, for the same time-constant given in the previous example, R will have to be smaller. This is due to the fact that, all other things being equal, the load capacitance will be greater because it now includes the capacitance of the pulse deck, its power isolation transformer, and whatever capacitance is associated with isolating the low-level drive signal, which now has a common-mode voltage equal to the pulse-output voltage. The common-mode voltage must be rejected.

There are two components of power dissipation in both the switch tube and the pull-down resistor. One is proportional only to the duty factor and the other only to repetition rate. The dissipation in the resistor during the pulse is the most obvious. It is V_0^2/R , where V_0 is the output-pulse amplitude. The average dissipation component is this peak power multiplied by the duty factor. The

corresponding component of switch-tube dissipation is the anode-cathode voltage drop—which is required to sustain the intrapulse resistor current—multiplied by the current. The average is this peak anode dissipation multiplied by the duty factor. At the beginning of each pulse the load capacitance C_L must be charged to a voltage V_0 , giving it an energy change of $1/2 C_L V_0^2$. An equal amount of energy is dissipated in the switch tube each pulse to accomplish it. The average-power dissipation component is the per-pulse energy multiplied by the pulse repetition rate. At the end of each pulse, the same amount of capacitor energy is dissipated in the resistor (regardless of its value) in the act of discharging it to its pre-pulse voltage level. The average-power component is the per-pulse capacitor energy multiplied by the repetition rate. We can see that both duty factor and PRF are independently stressful operating characteristics. Neither can be ignored.

The average current that flows through the modulator circuit, either from the main high-voltage power supply, or from the adjustable amplitude-control supply, is the sum of the duty-factor-proportional component, $V_0/R \times DF$, and the repetition-rate-proportional component, $V_0 \times C_L \times PRF$. Where fast rise-and-fall times, high-duty factor, and high PRF are required, the average current and dissipation can become prohibitively high, leading us to the next modulator topology.

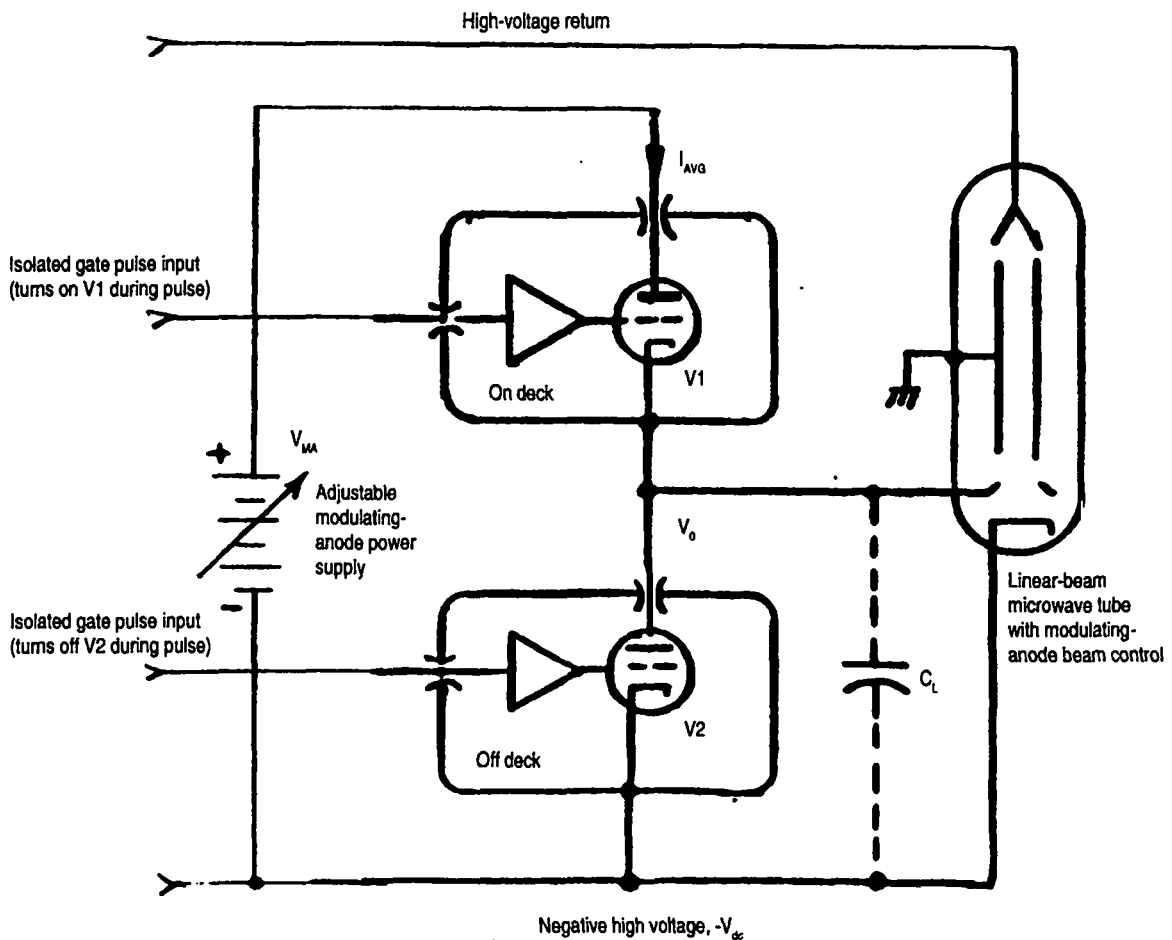


Figure 11-3. Modulating-anode pulser with active pull-up and active pull-down.

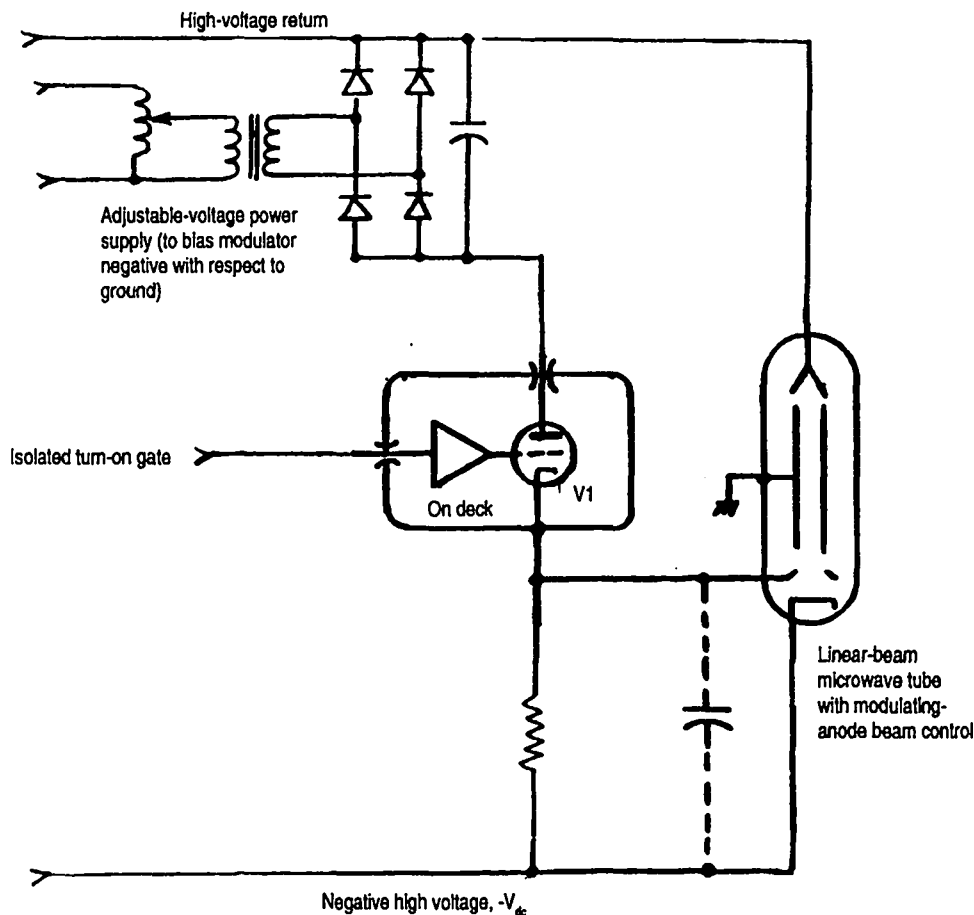


Figure 11-4. Alternative modulator amplitude control (why is this not likely to be a good idea?)

11.1.3 Active pull-up and pull-down

The topology of Fig. 11-3 has both active pull-up and pull-down, and the low-level gating strategy is such that the tubes never conduct simultaneously. This obviates a "shoot-through." If there is no resistive component to the load, such as modulating-anode interception current, there are no average current or dissipation components that are proportional to duty factor. All that is left are the PRF-proportional stresses. Both switch tubes experience average dissipation of $1/2 C_L V_0^2 \times PRF$, and the average modulator current is $V_0 C_L \times PRF$. The pulse rise-and-fall times depend entirely on the peak current available from the switch tubes, which, especially if they are tetrodes, will be nearly constant throughout most of the transition intervals until the diode-line voltages of the tubes are reached. After these voltages are reached, the current will fall off dramatically, and the time required for the last few percentages of voltage will be quite protracted. The initial rates of voltage rise-and-fall will be I/C_L , where I is the initial peak-anode current. Versatile, high-performance modulating-anode pulsers are almost always of this basic type.

However, if the type of power supply shown is used and variable amplitude control is desired (especially for high-ratio output, where pulse voltage is a large percentage of the load cathode voltage), then this topology can result in a discon-

certingly large amount of equipment, particularly if the modulator must operate at high PRF into large load capacitance. This situation raises the temptation to develop another circuit connection.

11.1.4 An attempt at simplified amplitude control

In the circuit of Fig. 11-4, the variable-amplitude power supply is now referenced to the positive return rather than the negative high voltage. If the modulator output is a large percentage of the beam voltage, the potential saving is obvious. Many modulators incorporating this strategy have actually been built. The problem is they often don't work, and the reason is obvious enough. The modulator current must pass through the amplitude-control power supply in the wrong direction and is, therefore, blocked by the rectifier diodes. In most such cases, however, the power supply has a bleeder resistor across it that modulator current will go through. If the resistor value is low enough so that the voltage produced across it by the average modulator current is not greater than the adjustable-power-supply output voltage, the circuit will work. When the pulse/voltage ratio is higher than 50%, it makes more sense to "buck," using a power supply referenced to the ground side, rather than "boost," using a power supply referenced to the negative high voltage, which is what the next circuit facilitates.




Figure 11-5. Another method of achieving modulator-amplitude control: the on-deck "grid-catcher" circuit.

11.1.5. The grid-catcher type of amplitude control

In Fig. 11-5 we see another active pull-up connection. Pull-down can be either active or passive, it doesn't matter. The "on" switch tube, $V1$, is cathode-driven by means of an FET in its cathode (as described earlier). The FET frees the grid of the tube for other purposes. In this case, grid voltage is returned to the deck potential through resistor R_G , which is shunted by capacitor C_G . The grid is also connected via a "grid-catcher" diode to the negative terminal of an adjustable dc power supply, whose positive terminal is connected to the positive return of the high-voltage power supply. During the interpulse interval, the FET is turned off, and the voltage at its drain rises so that the positive voltage at the cathode of $V1$ is sufficient to cut it off. At the beginning of a pulse, the FET is driven into conduction, which pulls the cathode of $V1$ near deck potential, which is also grid potential. The switch tube will conduct zero-bias anode current, charging up the load capacitance C_L at an initial rate of I/C_L , where I is the initial zero-bias anode current. When the output voltage attempts to become more positive than the voltage at the negative terminal of the amplitude-control power supply, the grid-catcher diode becomes forward-biased and begins to conduct, providing an alternative path for $V1$ output current. To pass through the diode, however, the current must also pass through R_G and C_G , charging up C_G in the direction that makes the grid of $V1$ negative with respect to its cathode, shutting it off again. (Note that this current is also in the proper direction through the amplitude-control power supply.) The output-pulse flat top will stop at a voltage with respect to ground that is more positive than the amplitude-control voltage by the amount of negative grid bias required to cut off $V1$ anode current. For instance, if the amplitude-control power-supply voltage was set at -10 kV and it takes a negative bias of 500 V to cut off $V1$, the modulating-anode pulse voltage will have a flat-top that is 9.5 kV negative with respect to ground.

During the flat-top portion of the output pulse, $V1$ functions as a cathode follower. Up to its limit, it will conduct whatever intrapulse current is required by the modulating anode or anodes. The only change in output voltage is the change in $V1$ grid-cathode voltage required for the current change. (The larger the $V1$ transconductance, the smaller the voltage change.)

If the grid-catcher diode were a perfect component with no shunt capacitance, there would be no need for the shunt capacitor C_G in the grid circuit of $V1$. Because there is capacitance across the diode, there will also be convection current through it during the rise-and-fall intervals of the pulse. This current would also flow through R_G , producing voltage across it that would tend to reduce the amount of anode current from $V1$. The negative voltage across R_G would be

$$\frac{dV_o}{dt} \times C_D \times R_G,$$

where C_D is the shunt-diode capacitance. Therefore, the value of R_G would, in effect, determine the rate-of-rise of the output pulse because it is I/C_L , where I is the $V1$ anode current.

The capacitor C_G integrates the dv/dt current, making it a smaller replica of the output-pulse rise-time wave shape. With a C_G large enough, the voltage

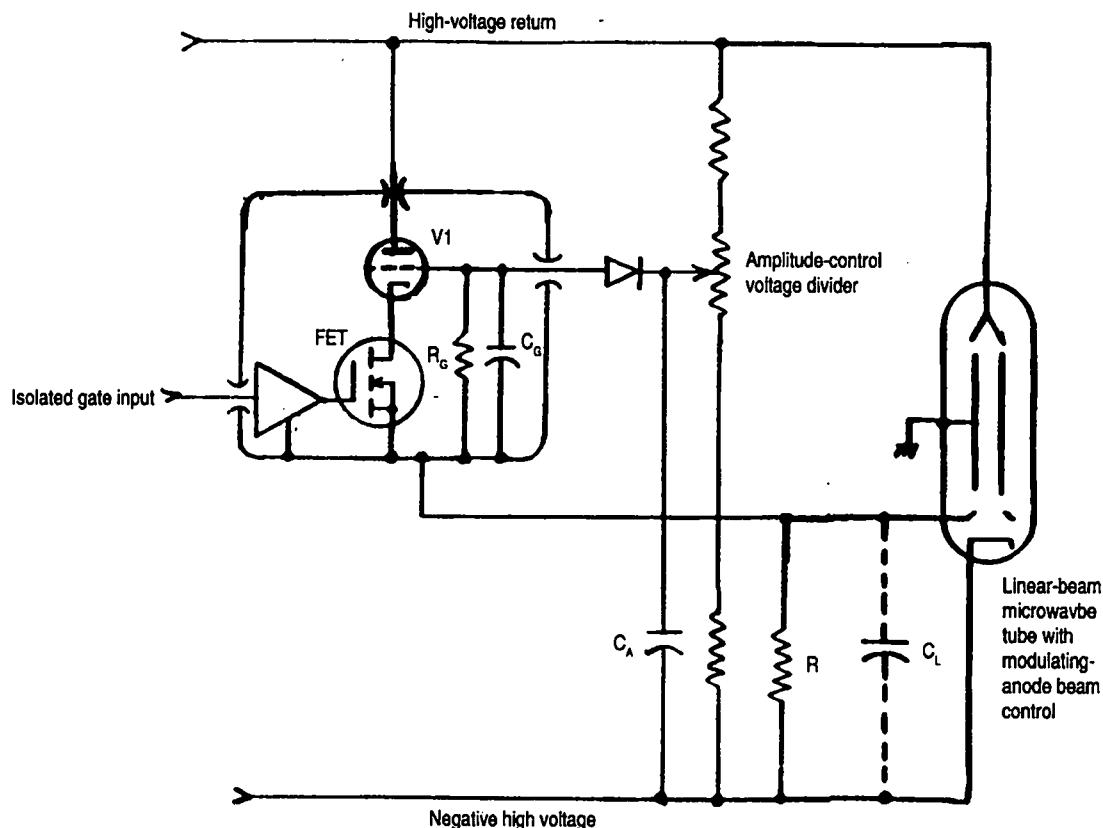


Figure 11-6. "Grid-catcher" type of modulator-amplitude control with passive voltage-divider adjustment.

across it can be made to have negligible effect on the $V1$ available current. The optimum value is usually experimentally determined, starting with one based on a reasonable estimate of diode capacitance—say, 5 pF—and reduced from that value successively until an effect on rise time can be discerned. If the switch tube is a triode, as shown, then note that its anode-grid capacitance will have the same effect as the diode capacitance and will have to be included with it. Both are magnified, in effect, by the "Miller" effect. The magnitude of C_G determines the value of R_G , which must be small enough so that $3 R_G C_G$ is less than the shortest interpulse interval. The value of these two components determines the average current that the amplitude-control power supply must handle. Once again, there are duty-factor and PRF-dependent components of the average current of the control power supply. The peak current through R_G is V_G/R_G , where V_G is the negative voltage required to cut off $V1$ anode current. The average-current component is this peak current multiplied by the duty factor. An amount of charge, $C_G V_G$, must be passed through C_G as each pulse top is established. The average-current component from this is $C_G V_G \times PRF$. The total control-power-supply average current is the sum of the two components.

11.1.6 The grid-catcher circuit with potentiometer-type control

Figure 11-6 shows an alternative to the amplitude-control power supply: an amplitude-control voltage divider. The potentiometer is obviously not just a simple, commercially available variable resistor. It is custom made. At least one

motor-driven version has been built using a multitude of small individual resistors in series with tap points between them. A screwlike mechanism driven by the motor moves the output lead from one tap point to the next in make-before-break fashion. The resistance of the voltage divider must be small enough so that the average diode current does not substantially alter the voltage at the tap point as duty-factor or PRF vary, as discussed above. The bypass capacitor C_A is made large enough so that the intrapulse value of diode current does not seriously alter the voltage across it. Note that this capacitor is referenced to the microwave-tube cathode voltage, which is subject to intrapulse droop as its beam current is removed from the main capacitor bank. This droop component is coupled by means of C_A to the amplitude-control tap point of the voltage divider. The modulating-anode pulse voltage, therefore, will track the negative high voltage, keeping the intrapulse beam current of the microwave-tube load essentially constant. (In fact, if C_A is not large enough, it will be charged up by the intrapulse diode current, and the modulating-anode pulse will actually rise with respect to the negative high voltage.)

11.1.7 *The gate-catcher type of amplitude control*

The circuits shown in Figs. 11-5 and 11-6 do not permit $V1$ to be operated with positive grid voltage. Often $V1$ is a tetrode, in which case adequate anode current is usually obtained with zero or slightly negative grid bias. But this is seldom the case with a power triode. Figure 11-7 shows how, with the addition of an additional FET stage in the $V1$ cathode circuit, any amount of effective positive grid drive can be introduced, as long as it is consistent with the drain voltage hold-off of the two FET switches. Even though the two FET switches are in series, they do not share the interpulse hold-off voltage at the cathode of $V1$ because they turn on and off at different times.

In the circuit illustrated, $Q3$ is the timing switch. Its gate drive responds to the optically coupled low-level turn-on gate at the input to fiber-optic-receiver, $U1$. The gate of $Q3$ is driven from an active pull-up/pull-down transistor pair, $Q1$ and $Q2$, that provides active source and sink for the dv/dt currents between the drain and gate of $Q3$. (These currents are akin to the "Miller" capacitance effect in a triode vacuum tube.) The $Q3$ circuit is a true bootstrap circuit, having its own isolated 15-Vdc power source, which is a high-isolation, low-input/output-capacitance dc-dc converter.

The other FET cathode switch, $Q4$, is already biased into conduction before a pulse begins. This is accomplished by the current flowing through $R1$ and $R3$ from a 24-Vdc source, which is also a high-isolation dc-dc converter. Zener diode $CR5$ clamps the $Q4$ gate voltage at +15 V. Even though $Q4$ is in the conducting state prior to the low-level input gate pulse, $Q3$ is still cut off, which keeps $V1$ cut off as well. The grid of $V1$ is "grounded" in that it is connected to the floating deck, from which the high-level pulse output is obtained. The cathode of $V1$ is not directly connected to the drain of $Q3$. (Because $V1$ in this case has a directly heated cathode, the cathode connection is made at the secondary center-tap of the filament transformer, $T1$, which is the electrical mid-point of the ac filament circuit.) Between the cathode and the $Q3$ drain is a dc power source comprising a bridge rectifier, $CR1-4$, and filter capacitor, $C1$. The ac input is

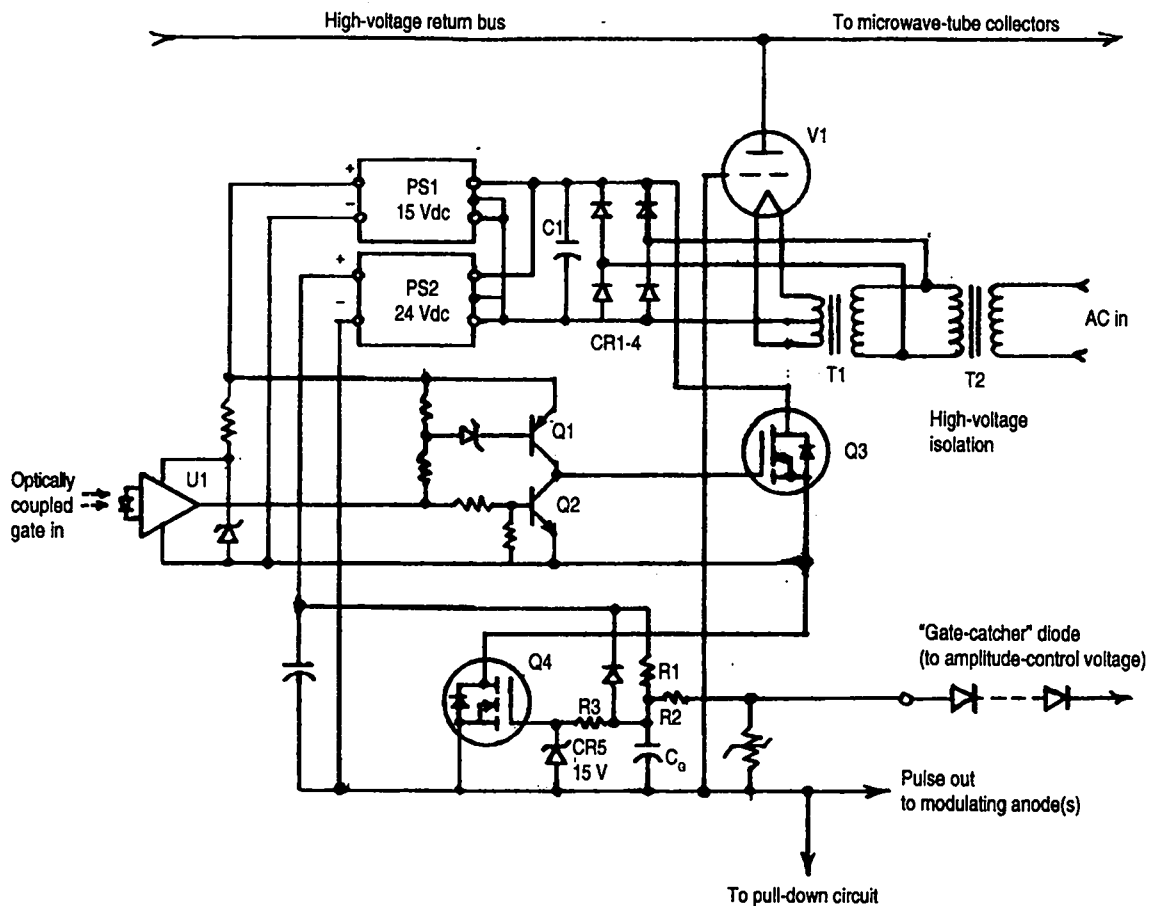


Figure 11-7. "Gate-catcher" circuit permitting switch-tube positive-grid control.

provided by the secondary winding of T2, the high-voltage isolation transformer. By coordinating with the primary voltage required of T1, the secondary voltage of T2 can be whatever is needed for the dc across C1 to be the desired amount of positive grid bias for V1. It is not applied to the grid, however. This voltage is applied to the cathode of V1 as negative voltage when Q3 is pulsed into conduction in response to the input timing gate. The entire cathode current of V1 must pass through this power supply, not just the grid current. During the interpulse interval, Q3 must hold off whatever cathode voltage is required to keep V1 cut off in addition to the bias voltage across C1.

Once Q3 is turned on at the beginning of an output pulse, conduction commences through V1, Q3, and Q4 (which was already on) to charge up the external load capacitance. When the output pulse, which increases from an interpulse value at or near the negative high voltage toward ground potential, reaches whatever voltage the anode of the "gate-catcher" diode string is connected to, an alternative current path through C_G and R2 becomes available. When this path is utilized, C_G rapidly discharges, shutting off Q4 and stopping the flow of load current, even though Q3 is still biased into its conducting state. The gate of Q4 now plays the same role in controlling output-pulse amplitude that the grid of V1 used to play.

Between pulses, C_G is recharged through R1, which is considerably smaller in

value than $R3$, to nearly 24 V, but the gate of $Q4$ is clamped to 15 V. Therefore, during the pulse rise time, the dV/dt current, which will flow in the stray capacitance shunting the control diode string, can discharge C_G by 9 V before any effect is noted in the state of conduction of $Q4$, thus permitting C_G to be made a smaller value than if it were directly connected to the $Q4$ gate.

11.1.8 The "quasi-resonant" modulating-anode pulser

So far, we have discussed circuit refinements that have reduced the operating stress levels to those proportional to load capacitance, pulse-repetition frequency, and the square of the output voltage. The product of these values will determine both the average-power dissipation in the modulator switches and the external average power required for the modulator to operate.

It is possible to contemplate a combination of parameters that results in power consumption and dissipation so high that the overall design becomes impractical, if not ridiculous. The topology shown in Fig. 11-8 may then become practical. Its obvious drawback is its complexity. It uses four electronic switches instead of two in the designs previously discussed. Its single, but perhaps crucial, advantage is that it provides true switch-mode application of the electronic switches. The load-capacitance charging- and discharging-current amplitudes are limited not by resistance, either internal or external to the switch tubes, but by reactance. This means that switching can be accomplished without power dissipation. During the leading edge of each pulse, both charge and energy are removed from capacitor $C2$, which is maintained at a voltage (with respect to the negative high voltage) that is one-half of the desired output-pulse amplitude.

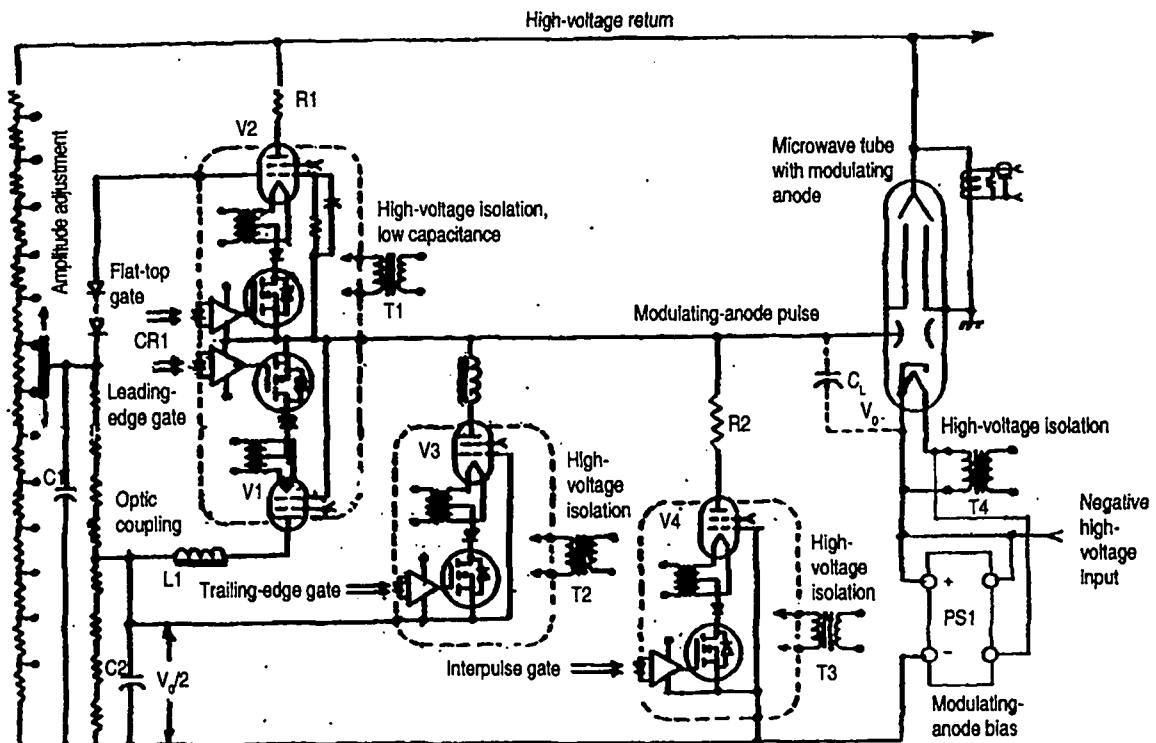


Figure 11-8. "Quasi-resonant" charge-and-discharge modulating-anode pulser.

The energy and charge are transferred to the load capacitance C_L through $L1$ and $V1$. Switch tube $V1$ is driven into a state of conduction such that its anode-voltage drop at the peak charging current is small with respect to the pulse output voltage. The smaller the switch-tube voltage drop, the less energy it will dissipate each pulse. Also the closer the charging-circuit impedance will approach

$$\sqrt{\frac{L1}{C_L}},$$

and the closer the peak current will approach

$$\frac{V_0}{2} \sqrt{\frac{L1}{C_L}}.$$

The output across C_L will be in the form of an inverted cosine wave for half a cycle, as shown earlier in Fig. 9-11. This wave described the dc-resonant charging of the capacitance of a pulse-forming network. The rise time, from zero to 100% of V_0 (which approaches 200% of the voltage across $C2$), is

$$\pi\sqrt{L1 \times C_L}.$$

The smaller the inductance of $L1$ can be made, the greater the peak current becomes—but the shorter will be the rise time. The current in the circuit is in the form of a half-sine wave. For this reason, even though $V1$ is shown in the form of a full-control electronic switch, it need not be. A half-control switch will be automatically commutated to the off state when the output voltage reaches 100% and the charging current attempts to reverse polarity. The switch (or a diode in series with it) must be capable of blocking a reverse voltage of $1/2 V_0$ as soon as the rise time is completed and the voltage across $L1$ collapses.

If we ignore for a moment the conditions during the flat-top portion of the output pulse, we can proceed directly to the trailing edge of the pulse, which is initiated by driving the discharge switch tube $V3$ into conduction in response to the optically coupled trailing-edge gate. The discharge path includes C_L , $V3$ (and its cathode-drive FET), and $L2$, whose rise-and-fall time performance is identical with $L1$. Current in this path restores to the source-energy store $C2$ the charge and energy that were removed from it during the rise time of the pulse. In fact, if the flat-top portion of the pulse were of zero length, the current through $C2$ and C_L would be a complete cycle of a sine wave, with the positive-going half flowing through $L1$ and $V1$ and the negative-going half flowing through $L2$ and $V3$. The output voltage V_0 across C_L would be in the form of a complete cycle of an inverted cosine wave, having a peak value of $V_0/2$, and, of course, a peak-to-peak value of V_0 , which is the output voltage we want.

Because we have produced sinusoidal waveforms by using two unidirectional switches conducting in time sequence, now we can delay the two half-cycles by

as much time as we choose (*sinusoidalus interruptus*), providing that we don't lose too much charge from C_L during the flat-topped portion of the pulse. This is the function of V_2 , which comes into conduction only when C_L has been fully charged. It acts as a trickle charger to provide whatever intrapulse leakage currents there might be, including actual beam-current interception by the microwave-tube modulating anode. (Even as little as 0.1% of a lot can be something.) Note that a grid-catcher diode, $CR1$, is used to coordinate the output voltage of the V_2 circuit with the "ring-up" voltage on C_L produced by the V_1 circuit. Capacitor $C1$ acts as the energy store for the grid-clamp diode current. Switch tube V_2 is the only one that must hold off the entire beam voltage of the microwave tube. The current it must handle, however, is a small fraction of what V_1 and V_3 must handle. For this reason, external resistance can be included in the anode circuit of V_2 that can be made high enough in value so that arc-through current can be made almost negligibly small.

The final switch tube, V_4 , conducts throughout the interval between the end of the discharge portion of one pulse and the beginning of the charge-up portion of the next. It provides a sink for whatever leakage currents there might be, including leakage across the ceramic insulator between modulating anode and body of the microwave tube. The leakage would otherwise tend to pull the modulating anode positive with respect to the cathode. (Note that a small positive interpulse modulating-anode voltage can be particularly devastating to a microwave tube because it gives rise to a minuscule amount of cathode current that is super-focused as it transits the length of the microwave tube. It then bores a very small hole right through the beam collector.) For this reason it is customary to provide a modest amount of modulating-anode bias voltage (1 kV to 2 kV) in order to assure that the modulating anode is negative with respect to the microwave-tube cathode throughout the entire interpulse interval. (This is the function of $PS1$.) The ac input to such a supply is often derived from the filament voltage of the microwave tube, as shown. Switch tube V_4 must be capable of holding off the full pulse voltage, V_0 , which is usually less than the full microwave-tube beam voltage. The current that it must handle is typically less than that of V_2 , so resistance $R2$ can be added to its anode circuit, serving the same purpose as $R1$.

It would be possible to configure the circuit so that a single inductor would suffice. It takes two inductors, however, to insure that all capacitive charging currents, including that of the floating-deck that is common to V_1 and V_2 , are limited by series-resonant impedance rather than by switch-tube internal resistance. Part of the total load capacitance that must be charged and discharged is the primary-to-secondary capacitance of the high-voltage power isolation transformer, $T1$. So it is important that this transformer be of low-capacitance design. The two other decks and the microwave-tube filament also require high-voltage isolation of the ac input power provided by $T2$, $T3$, and $T4$, but interwinding capacitance is of no particular importance, because the decks are referenced to nominally fixed voltages.

The modulator of Fig. 11-8, therefore, can actually be seen as a hybrid of two separate modulators. The first, using switch tubes V_2 and V_4 , can be seen to be a conventional active-pull-up, active-pull-down floating-deck pulser. Its role has




Figure 11-9. Hybrid modulating-anode pulser for velocity-jump guns (conditions during pulse).

been relegated to holding the flat-top of the output pulse up (V2) and the flat bottom of the interpulse interval down (V4). The other modulator is a quasi-resonant, or interrupted-resonance, capacitor-charging and -discharging circuit that comprises switch tubes V1 and V3 and resonating inductors L1 and L2, which operate from a secondary energy and charge store maintained at a voltage approximately 1/2 of the desired output pulse voltage. None of the switch tubes conduct simultaneously, so four separate, low-level timing signals are required, as shown. This is certainly not a simple modulator configuration, but it is able to satisfy otherwise unattainable performance requirements.

11.1.9 The modulating-anode power supply/cathode-pulser hybrid

Also in the category of hybrid is the arrangement shown in Fig. 11-9. It is best suited for microwave tubes that have "velocity-jump" electron guns and are characterized by a modulating-anode pulse voltage that is around half the cathode voltage. Extended-interaction oscillators and amplifiers that operate at moderate power levels are usually in this category.

In this case, there is no modulating-anode pulser as such. The microwave-tube modulating anode is connected to the negative terminal of a dc power supply, PS1, the positive terminal of which is connected to the collector (virtually at ground). During a pulse, the supply-output voltage continuously biases the modulating anode to the level that a conventional pulser would, except that its reference point is ground instead of cathode voltage. (This is really no different than the circuit of Fig. 11-5, except that there is no switch tube and no grid-catcher diode.) There is no beam current in the tube yet, because switch tube V2

and its cathode-drive FET block its path. The anode of the switch tube is connected to the cathode of the microwave tube through a coupling capacitor, $C1$. During the interpulse interval, the anode of $V2$ is maintained at a voltage close to ground by conduction through the "recharge" tube, $V1$, which can be replaced by a resistor if the operating duty factor is low enough. Diode $CR1$ keeps the cathode of the microwave tube positive by the amount of the diode drop (with respect to its modulating anode), which means that the average value of voltage across $C1$ is equal to the modulating-anode voltage as well.

To initiate a beam-current pulse, switch-tube $V2$ is driven into conduction in response to the optically coupled "on" gate, while $V1$ is driven from the conducting to the non-conducting state. The cathode-modulator dc power supply $PS2$ has an output voltage slightly larger (by the amount of the conduction anode-voltage drop of $V2$) than the voltage by which the cathode of the microwave tube must be pulled negative (with respect to its modulating anode) for the desired amount of beam current. For a 50%-modulating-anode tube, this means that $PS1$ and $PS2$ have very nearly the same voltage ratings, each one-half of the microwave-tube beam voltage. When $V2$ conducts, it pulls the positive end of $C1$ negative by an amount slightly less than the $PS2$ voltage. Assuming that $C1$ has a sufficiently large capacitance, the negative end follows the positive end, and the cathode of the microwave tube, which was already negative by an amount equal to the $PS1$ voltage, is pulled even further negative so that the intrapulse value is very nearly $PS1 + PS2$ (assuming zero conducting-anode drop for $V2$).

The path of pulse current is shown by the arrows. It includes $C2$, the microwave-tube collector-cathode path, $C1$, and $V2$. Figure 11-10 shows what happens between pulses. During the pulse, charge is removed from both $C1$ and $C2$ that is equal to

$$\int I \cdot dt,$$

where I is the peak-beam current and dt is the pulse duration. Except in long-pulse, high-duty-factor situations, there is negligible contribution to pulse current directly from the dc power supplies. Current from them restores the lost capacitor charge during the interpulse interval. There are two separate current components: I_1 is current from $PS2$ restoring the lost charge to $C2$; I_2 flows out of $PS1$ through the recharge tube $V1$, which conducts during the interpulse interval, and through $C1$ and $CR1$, restoring the lost charge to $C1$. The average values of the two currents are the same and equal to the time-averaged value of microwave-tube beam current. For some people, the only surprise of this arrangement is that the average current from a power supply connected to a modulating anode, which itself draws no current, is equal to the average beam current. In this case, the total power requirements of a 50%-modulating-anode-voltage tube are shared almost equally by the two dc power supplies, both in voltage and current. Regardless of the voltage split between modulating-anode voltage and cathode voltage (or the degree of "voltage jump"), the average currents from the two power supplies will be identical.

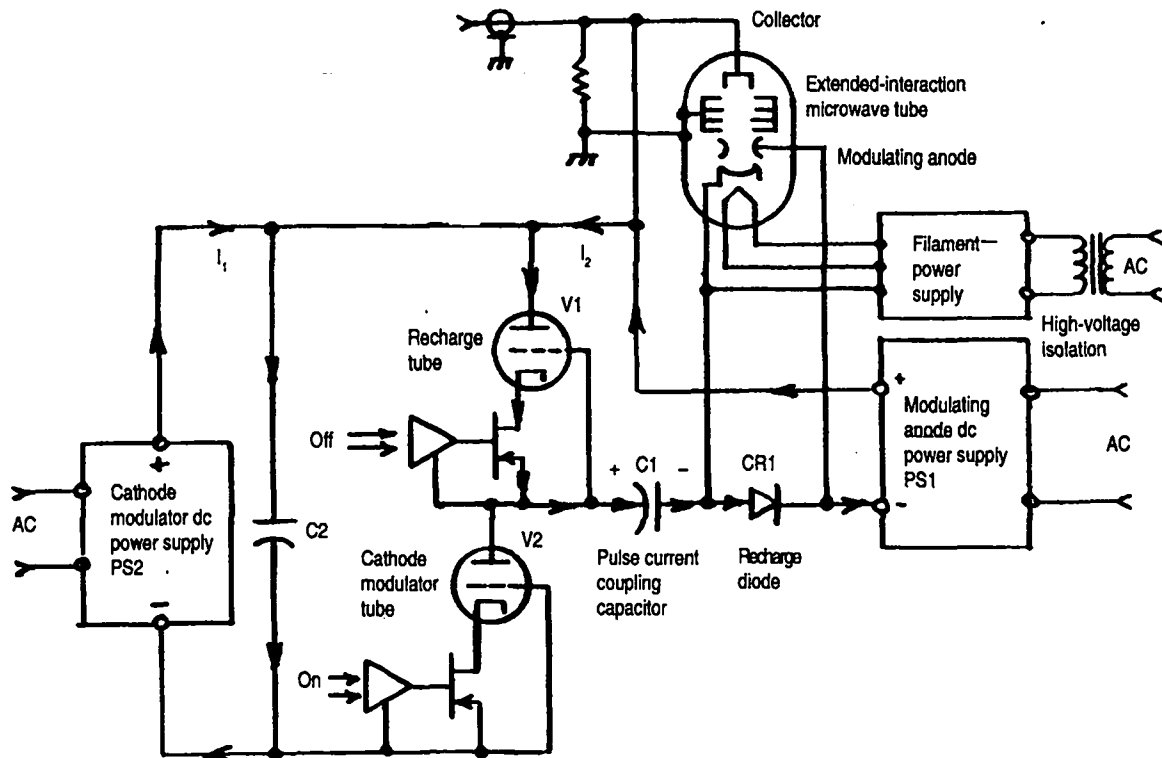


Figure 11-10. Hybrid modulating-anode pulser for velocity-jump guns (conditions between pulses).

11.2 The self-capacitance of modulator decks

If a significant portion of the load capacitance that a modulating-anode pulser must drive can be embodied in the electronic equipment deck associated with the pull-up vacuum tube, it stands to reason that a good design objective would be to minimize it. There is only one strategy that will do this: making the deck as physically small as possible and spacing it as far as possible from everything else, until the limiting case of corona-discharge threshold is reached. At this point, the deck will be a sphere of radius r . If V is the deck voltage, V/r must be less than the local-breakdown electric-field intensity for the medium in which it is immersed. (If, for instance, the medium is air, at standard temperature and pressure, V/r must be less than 30 kV/cm). However, practical considerations usually result in deck designs having dimensions that are considerably greater than the limiting case. (In these instances, the form factor looks more like a rectangular box with rounded corners than a sphere.)

Figure 11-11 is intended to provide an intuitive, as well as numerical, assessment of the difference in capacitance between a square-sided deck and the remotely located walls of a larger containment enclosure. The piece-wise approximation is based on the equivalent series capacitance of successive virtual-parallel-plate capacitors created by slicing up the space that extends outward from the deck surfaces, as shown by the dashed lines. Only one of the six contiguous volumes is shown. (Common surfaces are indicated by 45° angles.) If the length of each side is L , and we choose to slice up the space that progresses outward from each side into slabs of incremental thickness $0.1L$, then it can be seen that the dimensions of each successive slab are $L(1 + 0.2n)$ on a side, where n is the




Figure 11-11. A method for determining modulator-deck self-capacitance.

number of the particular slab, with 1 being the closest one. Therefore, the area (A) of each slab is $L^2(1 + 0.2n)^2$. Except for their ends, the dashed lines dividing up the slabs lie on electric-field equipotentials, so they could be replaced with conducting surfaces without altering the electric-field distribution. The parallel-plate capacitance between surfaces C_n is given as

$$\frac{0.225A_n \text{ pF/in.}}{0.1L}$$

normalized to the length of a side, L . The area, and therefore the parallel-plate capacitance, of each successive slab is greater than its immediate predecessor that is closer to the deck. The deck's total capacitance is the equivalent series capacitance of all of the incremental slabs, or

$$\text{Capacitance / side} = \frac{1}{\frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} + \dots + \frac{1}{C_n}}$$

and

$$\text{Total capacitance} = 6(\text{capitance / side}).$$

Table 11-1 shows how the areas and capacitances of successive slabs increase,

n	Area/slice (in. ²)	Capacitance/slice (pF)	Total Capacitance (pF)
1	1.41 L ²	3.17 L	3.17 L
2	1.96 L ²	4.41 L	1.84 L
3	2.56 L ²	5.76 L	1.4 L
4	3.24 L ²	7.29 L	1.17 L
5	4.0 L ²	9.0 L	0.78 L
6	4.84 L ²	10.9 L	0.73 L
7	5.76 L ²	13 L	0.69 L
8	6.76 L ²	15.2 L	0.66 L
9	7.84 L ²	17.6 L	0.64 L
10	9.0 L ²	20.3 L	0.62 L
11	10.2 L ²	23 L	0.60 L
12	11.6 L ²	26 L	0.59 L
13	13.0 L ²	30 L	0.57 L

Table 11-1. How capacitance increases with slab area while series-equivalent total capacitance decreases with each additional slab.

and how the series-equivalent total capacitance decreases with each additional slab. But note that the total capacitance does not decrease very rapidly after about the eighth or ninth slab. This means if an outer wall is farther away than 0.8 or 0.9 times the deck length, it doesn't matter how far away it is—or what shape it is, either. The curve shows that capacitance per side is asymptotic to approximately $0.6L$ pF. (This calculation assumes there is a relative dielectric constant of unity for the immersion medium, which certainly isn't true if the medium is dielectric oil.) An air-insulated deck, with dimensions of 6 in. on a side and with containment walls at least 5 in. away, can be expected to have a capacitance per side of approximately 0.6×6 in., or 3.6 pF, and a total capacitance of approximately 6×3.6 pF, or 22 pF. The same deck immersed in transformer oil that has a relative dielectric constant of 2 will have twice the self-capacitance. Increasing the size of the containment vessel will have virtually no effect on the deck self-capacitance, whereas changing the dimensions of the deck sides will have a direct effect on it.

Given that the deck will have to support at a minimum the modulator pull-up vacuum tube, the footprint of its tube socket will establish the minimum area for two of the deck sides. The largest single component that must be housed within the deck is usually the filament transformer for the switch tube. For this reason, the secondary voltage of the low-capacitance, high-voltage isolation transformer is sometimes made equal to the required filament voltage, thus eliminating the need for the filament transformer altogether. This, however, requires that other transformers within the deck be specially wound so that their primary voltages are also the same as the filament voltage. Transformer bulk can also be reduced by using higher-frequency primary power, such as the standard 400 Hz. Solid-state static inverters that operate at this frequency are readily available, as are many transformers of different designs. Even higher frequencies than 400 Hz are practical, but the advantages gained are seldom worth the high cost of the special-purpose components designed to exploit them.

11.3 Vacuum tubes appropriate for modulating-anode-pulsar service

The obvious performance criterion of a switch tube for modulating-anode-pulsar service is an adequate voltage hold-off capability. This, in turn, depends upon the circuit topology chosen (or invented) and the modulating-anode-voltage excursion required by the microwave tube, or family of such tubes involved. In reviewing the topologies previously discussed, we see that required voltage hold-off can range from one-half of the required voltage swing (the leading-edge and trailing-edge stages of a quasi-resonant modulator) to an amount equal to the full cathode voltage of the microwave tube (the pull-up tube of grid-catcher type connection).

The amount of peak-pulsed current required depends only upon the total load capacitance and the desired charge-and-discharge time intervals (rise time and fall time). The charge transport is the same for each pulse, so average current is determined by voltage swing, load capacitance, and pulse-repetition rate (unless the resistive pull-up or pull-down strategy is used). Average current, therefore, is independent of transmitter duty factor as such. Consequently, the switch-tube duty factor is usually a small fraction of transmitter duty factor. Tubes appropriate for this service, therefore, tend to be considerably smaller—electrically and physically—than those required for use as hard-tube modulators, such as beam-pulsars. Figure 11-12 shows the relative envelope sizes, compared with a 12-in. reference line, of some tube types that have been used in successful modulating-anode pulsars, even though they were not necessarily designed for such service specifically. (The tubes discussed are all members of families of similar tube types and are by no means the only ones worthy of consideration when contemplating a new application.)

11.3.1 *The 4PR250C/8248 tetrode*

Labeled A in Fig. 11-12, this glass-envelope, radial-beam tetrode switch tube has seemingly been around forever and has graced the sockets of a multitude of air-insulated modulating-anode pulsars. Its anode-cathode voltage-hold-off capability is a very honest 50 kV, and although it is not rated for higher voltages, it has been known to tolerate more. Its long-life thoriated-tungsten filament will produce useful plate current greater than 4 A, and its radiation-cooled tantalum anode will dissipate 250 W. It is also relatively affordable, making it an excellent choice in pulsar designs operating up to 50 kV.

11.3.2 *The 8960 tetrode*

The 8960 tetrode (B) is another glass-envelope, radial-beam tube with an honest 50-kV anode-voltage hold-off capability and a thoriated-tungsten, directly heated cathode. But it is physically and electrically larger than the 4PR250. It boasts 12-A peak-anode-current capability and 1200-W anode dissipation. It is often employed in air-insulated modulator designs up to 50 kV that are beyond the capability of the 4PR250.

11.3.2 *The Y-847 triode*

A relatively modern tube, the Y-847 triode (C) has planar-electrode geometry and ceramic-metal construction that is especially designed for high-voltage service. Its construction makes it independent of orientation and tolerant of high




Figure 11-12. Tube types well suited for modulating-anode pulser service.

levels of vibration and shock. Unlike the previous tubes, which were much more mature, it has a unipotential, oxide-coated, indirectly heated cathode that is designed for short-pulse (20- μ s) service. Its peak-pulse cathode current is rated at 10 A. The relatively short pulse-duration rating is not a serious limitation in modulating-anode-pulser service because switch-tube conduction is required only during either the rise or fall interval, anyway. Its amplification factor, μ_A , is 1300, which gives a pretty good idea of how close the grid is to the cathode surface in relation to the cathode-anode spacing. This geometry not only contributes to the dc anode-voltage hold-off rating of 150 kV but to a very short grid base as well (~100 V grid-cathode for plate current cut-off at 150-kV anode voltage). The triode is intended for use in oil-insulated applications, and its anode can dissipate 2000 W when an external radiator is bolted on. (The radiator is shown in the figure.)

11.3.4 YU-146 tetrode

The tube shown as D in Fig. 11-12 is another member of the growing family of modern ceramic-metal, high-voltage switch tubes, except that it is a tetrode with planar geometry. Sharing the ruggedness of the YU-847 triode, the YU-146 is capable of sustaining 60-G shock and 10-G vibration. It is rated for 175-kVdc anode hold-off voltage when oil-immersed with proper corona shielding (not shown in the figure). Its peak-pulse cathode current, from its unipotential, dis-

penser-type cathode, is rated at 20 A for 11- μ s-duration pulses. Its amplification factors, μ_a and μ_{sg} , are approximately 2000 and 150, respectively, and its transconductance, which is related to perveance, is 0.15 mho ($\Delta I_p/\Delta V_G$). The anode dissipation is dependent upon its immersion medium, tube orientation, and whether convection is forced or natural. With forced oil convection and horizontal mounting, dissipation can be as high as 10,000 W.

11.3.5 TH-5184 tetrode

The TH-5184 tetrode (E) is a tube designed for use in pulsed x-ray equipment. Its anode-voltage hold-off is rated at 85 kV and anode dissipation at 1000 W when immersed in oil. These parameters also make it a serious candidate for modulating-anode-pulser service. Its design is more conventional. It is a glass-envelope tube with a thoriated-tungsten cathode capable of peak-pulse cathode current of 5 A. Also more conventional are its amplification factors, with a μ_A of approximately 160 and μ_{sg} of approximately 6. Its transconductance is 0.01 mho, which is small compared with the much higher value for the more modern planar-geometry YU-146 tetrode.

11.3.3 TH-5188 tetrode

The tetrode illustrated as "F" in Fig. 11-12 is the largest—electrically and physically—and most recent member of this tube family. The TH-5188 is rated to 120-kV anode-voltage hold-off and up to 10,000 W of anode dissipation when operated in oil. Its larger thoriated-tungsten cathode is rated for peak-pulse cathode current of 10 A. Its amplification factors are approximately 500 for μ_A and approximately 8 for μ_{SG} . Its transconductance is 0.02 mho.

11.4 Some representative modulating- anode pulser designs

11.4.1 The BMEWS dual-klystron modulator

The dual-klystron transmitters designed for use in the Ballistic Missile Early Warning System (BMEWS) used the very first radar-type modulating-anode klystron, the X-626, which operated in the UHF band. It was rated at a peak RF power output of 1.25 MW when operated with 2-ms-duration pulses at a 30 pps repetition rate, with continuously applied cathode voltage of up to 120 kVdc, and with modulating-anode voltage swing of about 2/3 of the cathode voltage. The physical arrangement of the dual klystrons and their floating-deck modulating-anode pulser, which may very well have been the first of its type, is shown in Fig. 11-13.

The transmitter was designed by Continental Electronics, then and now a designer and manufacturer of high-power radio-broadcast transmitters. It was built at a time when fiber-optic links were unheard of. It should not be surprising to learn that the then-very-daunting problem of small-signal coupling to the two floating electronic decks, especially the "on" deck, was solved by radio broadcast. Each of the two pulse decks, "on" and "off," had its own receiving loop antenna. Between them was a third loop antenna, the transmitting antenna. For the duration of each desired transmitter output pulse, a burst of a pulse-modulated, 5-MHz, radio-frequency signal was fed into the transmitting loop from a not-inconsiderable RF source. The signal was simultaneously picked up by the

two receiving loops and rectified within the two decks to produce a gating signal to the high-voltage switch tubes. (This topology is similar to the one shown in Fig. 11-3.) The gate turned off the "off" switch tube and turned on the "on" switch tube. The anode of the "on" switch tube was maintained at a positive voltage (with respect to klystron cathode voltage) of approximately two-thirds of its cathode voltage by means of a fixed resistive voltage divider and shunt energy-storage capacitor. The resistive/capacitive voltage divider supplanted the function of the variable power supply shown in Fig. 11-3. It was possible to do this because the repetition rate (30 pps) was invariant and, therefore, the average modulator current was constant.

This class of transmitter, which became populated by ever-more-modern klystrons over its long life span, has only recently been replaced by highly modular, solid-state phased-array radar systems of the Pave Paws type. The decommissioned klystron transmitters are usually snapped up eagerly by other users, by the way.

11.4.2 The Millstone Hill radar transmitters

One such eager user of BMEWS-class transmitters was the Millstone Hill Radar Facility at Westford, Mass., which employs three of the dual-klystron transmitters. Two use UHF klystrons. The third transmitter uses X-780 L-band




Figure 11-13. An early (if not the first) modulating-anode pulser.

klystrons, which produce 2.5 MW peak power each. Of the UHF transmitters, the first started with the ancient X-626 external-cavity klystrons, which were later replaced by the L-3403. At present, one of the two transmitters has been upgraded to the high-efficiency version of the L-5773, which produces almost as much power from the two klystrons as was previously obtained from all four combined.

Originally, the modulators for these transmitters were the BMEWS type. They had virtually no repetition-rate agility or amplitude variability, and their rise-and-fall times were excessively long for the radar experiments that were later




Figure 11-14. The modulating-anode pulser for the Millstone Hill dual-klystron radar transmitter.

contemplated. For these reasons they were eventually modified to the type shown schematically in Fig. 11-14. This design represents a classic embodiment of the active-pull-up, active-pull-down, grid-catcher amplitude-control topology using TH-5188 tetrode switch tubes. Even though the basics of this arrangement have been previously discussed, there are a few changes worth noting.

The first is the use of an external resistance in series with each anode of the two switch tubes. This resistance serves dual purposes. It is small enough (3 kohms each) so that it does not significantly reduce peak-pulse switch-tube anode current, which means that it is not large when compared with the incremental anode resistance of the tubes. It will, nevertheless, share the total power dissipation incurred in repetitively charging and discharging the load capacitance (some 600 pF) with the switch-tube anodes. More importantly, it limits the "shoot-through" fault current that may result from simultaneous arcs in the "on" switch tube and in the klystron between the modulating anode and cathode, or simultaneous arcs in the "on" and "off" switch tubes. Neither coincidence is impossible.

Note also that the output from the "on" deck and the return connection from the "off" deck are diode-coupled from the screen grids of the switch tubes rather than from the deck connections. This is because an internal arc in either switch tube will be from anode to screen grid. The fault current, therefore, is diverted to the external circuitry by way of the diodes, thus minimizing the likelihood of arc-through to control grid and then cathode, an event that could greatly jeopardize the switching transistors. For the "off" deck, this approach provides the additional advantage that the screen-grid power supply biases the source electrodes of the switching transistors negative (with respect to the klystron cathode voltage). These transistors are in the conducting state during the interpulse interval, turning on the "off" switch tube. Once saturation conduction has been achieved, its anode is therefore also negative (with respect to the klystron cathode voltage), providing negative modulating-anode bias for the klystrons during the interpulse interval. In this circuit there is an additional bias power supply. (Actually, the screen voltage is taken from the midpoint of a single power supply.) The price for double duty from the screen supplies is that they must be rated for the entire average current of the modulator rather than just for the average screen-grid current.

The two switching transistors in the cathode circuit of the "off" tube, *Q1* and *Q2*, provide two states of conduction for the "off" switch tube. At the beginning of the fall time of each pulse, both transistors are pulsed into conduction. Transistor *Q2*, which has no resistor in its drain circuit, conducts the full discharge current of the "off" deck, rapidly terminating the pulse. After a time delay of slightly longer than the normal fall time, *Q2* is gated off. Transistor *Q1*, however, continues to conduct. The resistor in its drain circuit is selected to cathode-bias the switch tube so that plate current is a few milliamperes but the screen-grid current is negligible. The tiny bleed current is usually adequate to handle leakage current that may flow across the inner surfaces of the ceramic insulators between the modulating anodes and bodies of the klystrons. (This leakage is due to barium-based deposits from the cathode.) Without a current sink with low voltage drop, which the barely conducting switch tube provides, the leakage

current could pull up the klystron modulating anodes sufficiently to cause low-level, but continuous, beam current, which can be extremely injurious to the klystrons. (Actually, large klystrons with oxide-coated cathodes are less likely to exhibit ceramic-bushing-leakage current than some of the more modern, high-power travelling-wave tubes that have dispenser-type cathodes.)

The new modulators at Millstone Hill were successful. They have permitted transmitter operation at pulse repetition rates up to 1000 pps with pulse rise-and-fall intervals on the order of 10 μ s. And they have allowed the use of complicated waveforms.

11.4.3 *Dual-klystron modulating-anode pulser for particle-accelerator RF source*

Not all modulating-anode pulsers are built for dual-klystron transmitters, but the example shown in Fig. 11-15 is another one. In this case, the microwave amplifier tubes are 1.25-MW-peak-power klystrons, operating at 850-MHz and a pulse-repletion rate of 10 Hz. They require approximately 90 kVdc for their cathode voltage and are used as the input to drift-tube linear-accelerator cavities. What is unique about this modulator, which has an active-pull-up, passive-pull-down topology, is that it has a two-stage grid-catcher circuit for its pull-up pulse deck.

The high-voltage electronic switch is an L-5012 beam-switch tube that is used at only a fraction of its collector-current capability. Resistive pull-down is tolerable because the pulse duration is long, 2 ms, and the maximum accelerator duty factor is only 2%. Furthermore, this application of RF power is unlike a monostatic radar-system transmitter, where a klystron's beam-induced receiver noise following the transmitter pulse can limit the system's minimum range. Such noise is not a concern in an RF source for a particle accelerator. The 100-kohm pull-down resistor, *R1*, produces a sufficient fall time.

The modulating anode of the L-5012, *V2*, is pulsed to approximately 2.5 kV positive (with respect to its cathode) by a bootstrap-connected driver stage (a bootstrap within a bootstrap) comprising a 4PR125 tetrode, *V1*, and a cascode-connected FET switch, *Q1*. The return for *Q1* and the optically coupled low-level gate-drive circuit is the modulating-anode of *V2*. The grid-catcher diode *CR1* is connected to the grid of *V1*. When the optically coupled "on" gate is received, *Q1* is turned on. This transistor pulls the cathode of *V1* low, pulsing it into a zero-bias conduction state. This action connects the positive terminal of the 2.5-kVdc power supply to the modulating anode of *V2*, pulsing the L-5012 into conduction, which commences charging the load capacitance and building up current in *R1*. This build-up continues until the modulator output voltage is approximately 2.5-kV more negative than the voltage across *C1*, which is the storage capacitor for the voltage-divider-type amplitude control (as in Fig. 11-6). When this point is reached, *CR1* begins to be forward-biased and the shut-off process of *V1* commences, clamping the modulator output voltage.

When the low-level gate signal ends, drive is removed from the modulating anode of *V2* and the trailing-edge commences, the fall time of which is primarily determined by the value of *R1*. The value of the resistance shunting the *V2* modulating anode to its cathode circuit can also affect fall time because it is the sink for the "Miller" capacitance current that flows between the *V2* collector and

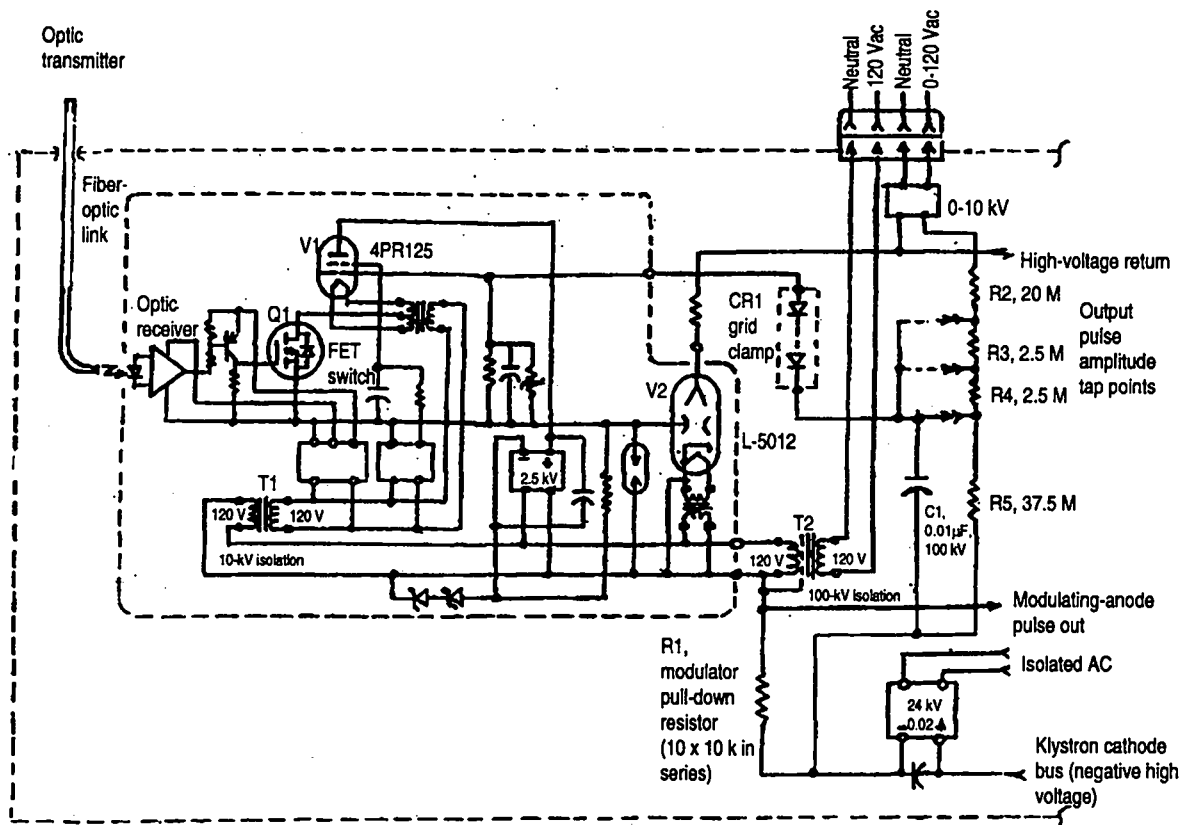


Figure 11-15. Simplified schematic diagram of dual-klystron modulating-anode pulser for particle-accelerator high-power RF system.

modulating anode.

Exploded views of the physical arrangement of the dual-klystron RF station are shown in Fig. 11-16. The self-capacitance of the modulator pulse deck is quite high because the dimensions of the deck are not only relatively large, 10 in. x 20 in. x 17 in., but the deck is immersed in a tank of synthetic dielectric oil having a relative permittivity of 3.2. If we calculate the self-capacitance of the deck in air using the methods described in Section 11.2, a value of approximately 120 pF is obtained. A measurement of the total capacitance from modulating-anode bus to ground, with no oil in the tank, gives a value of approximately 230 pF. This includes the modulating-anode capacitance of the two klystrons (a total of about 100 pF) and the capacitance of the power isolation transformer for the pulse deck (about 11-15 pF). All of the capacitance, except that inside the vacuum envelopes of the klystrons, will be multiplied by a factor of 3.2 when the tank is filled with dielectric fluid, bringing the total capacitance to about 500-600 pF. The discharge time-constant of the passive pull-down is, therefore, approximately 50-60 μ s. But the measured trailing-edge discharge time-constant operating in oil is 80 μ s.

This additional time is introduced by the "Miller-effect" time-constant of the switch-tube driver. These two time-constants, which are isolated from one another but are cumulative, are added together as the root of the sum of the squares of the individual time-constants. The resulting calculation of the driver-stage

discharge time-constant indicates it is approximately the same as the external pull-down constant. (If the time constants were exactly equal, they would produce a result that is the square root of 2 multiplied by each individual time-constant.)

A note of caution: measuring low-level signals within pulse decks and high-voltage, or "hot," decks can be lethal. To make such measurements requires that an oscilloscope or meter be "floated" so that its signal common is at deck potential. Failure to observe this rule can have a tragic outcome. (At least one engineer and one technician have been killed as a result of such carelessness. The temptation to reach out and make adjustments to the test equipment is often irresistible—but never habit forming.)

11.4.4 The modulating-anode pulsers for the MAR-1 radar system

The Multi-function Array Radar 1, or MAR-1, was developed during the mid-1960s. It was an experimental, phased-array (PA) radar designed to be part of the Safeguard anti-ballistic-missile defense system, which was the first system to time-sequentially accomplish the functions of target acquisition, tracking, and discrimination within a single pulse-repetition interval. It achieved this by using the electronic agility and broad bandwidth of true time-delay steering, which governed the separate transmitting and receiving antenna arrays. This radar was the test bed that led to powerful PA systems such as Cobra-Dane, Cobra-Judy, the Perimeter-Acquisition Radar, and even the SPY-1 radar of the Aegis weapons system. Most of these systems are still operational. But unlike its successors, MAR-1 was purely experimental. It had no long-term operational mission. (Part of the experiment was to see if it could be operated by Army enlisted personnel.)




Figure 11-16. Physical arrangement of 850-MHz dual-klystron RF station.




Figure 11-17. Simplified representation of the MAR-1 transmitter-array face.

It could.)

Its transmitting-array antenna, which used a total of 824 TWT RF power amplifiers, was electronically steered by the switching of incremental time delays that took the form of transmission-line segments, which were graded in delay time in accordance with a binary progression in the low-level input paths of each of the TWT PAs. A feature of this transmitter, shared by none other, was the means of focusing the 824 TWTs, which is illustrated in Fig. 11-17. Whereas tubes of this type, which have 5 kW of peak-power output at L-band, are routinely focused by means of periodic permanent-magnet stacks or solenoidal electromagnets that surround their 3-ft-long interaction regions, none had ever been focused by a uniform-field permanent magnet system. The reason for this is the field produced by a permanent magnet for a tube that is long and skinny is inefficient. Most of the magnetic energy will be located in the fringing fields.

If, however, one approaches the focusing of all 824 TWTs as an ensemble, what can result is what is shown in Fig. 11-17: a single, monolithic, permanent magnet whose iron pole-pieces are 30 ft on a side but separated by only 3 ft. The form factor of the design has gone from long and skinny to short and fat, and for that reason the magnetic energy lost to fringing fields becomes a trivial portion of the whole. Just such a magnet was built and energized by a total of 1000 permanent-magnet rods made of Alnico 5. The individual TWTs, for their part, had dislike iron field-straighteners arrayed along their lengths to short-circuit transverse magnetic-field components along their beam axes. The combination worked.

Although one of the candidate TWTs had a gridded-electron gun, it was

Group	TWT count	Group	TWT count
A	41	K	41
B	41	L	41
C	41	M	41
D	41	N	41
E	41	O	41
F	41	P	41
G	41	Q	42
H	41	R	41
I	41	S	42
J	40	T	40

Table 11-2. Grouping of the 824 TWTs for the MAR-1 transmitter.

decided that operation of multiple tubes from a single bus would be more feasible if the tubes had full-voltage modulating anodes instead. For the purposes of dc and modulation-pulse inputs, the 824 TWTs, were divided into 20 separate groups, A through T, as shown in Table 11-2. Each group had as many as 42 TWTs as loads, and each group had its own dc power supply, energy-storage capacitor bank, operating controls, and, of course, modulating-anode pulser. Unlike most such applications, these TWTs were physically remote from their modulating-anode pulsers and were interconnected by high-voltage coaxial pulse cables having significant distributed capacitance (at a rate of typically 30 pF/ft). Furthermore, the TWTs associated with a given pulse modulator were not placed in adjacent rows; the original design called for three separate pulse cables from each modulator to the feed ends of three cable trays. (Before the transmitter became operational, however, this design was modified. A single cable from each modu-

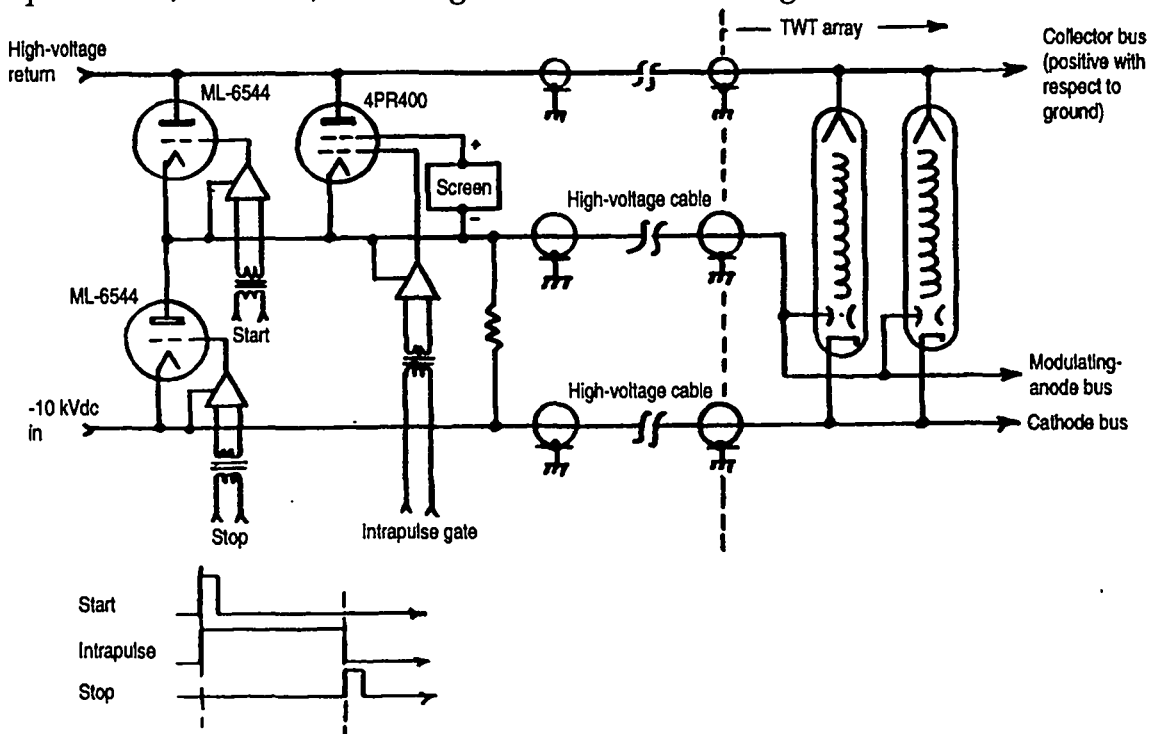


Figure 11-18. Simplified schematic diagram of one of twenty MAR-1 transmitter pulse modulators.

lator was used. This cable, in turn, fed into an intermediate junction box at the periphery of the magnet assembly, from which shorter jumper cables split out to the individual cable trays. This also solved the problem of a 2-MHz oscillation that was superimposed on the modulating-anode output voltage.)

The pulse modulators themselves, shown in Fig. 11-18, had active pull-up and pull-down provided by two ML-6544 triodes, which are more commonly encountered in cathode pulsers. These tubes conducted only during the rise-and-fall intervals of the output pulse, charging and discharging the load capacitance and making full use of the short-pulse, high current available from their oxide-type cathodes. A third switch tube, a 4PR400, had a long-pulse thoriated-tungsten cathode. It is shown shunting the pull-up tube. Its function was to supply the significant intrapulse modulating-anode interception current for as many as 42 parallel TWTs. Even if individual TWT modulating-anode current was 0.1% of cathode current, the total current for 42 tubes is 4.2% of total cathode current, an amount that cannot be ignored.

The start, stop, and intrapulse low-level gate signals were coupled to the switch-tube drive circuits by means of special pulse transformers that had a high degree of electrostatic isolation. Even so, the common-mode-rejection problem was formidable, especially for the pull-up deck. And this problem was further aggravated by the use of SCRs to discharge small PFNs, which functioned as the grid-drive-circuit active elements. The SCR gate circuits proved to be far too sensitive for such an application, although they were eventually made to work.

11.4.5 The long-range imaging radar transmitter modulating-anode pulsers

Located atop Haystack Hill (the next hill north of Millstone Hill near Westford, Mass.), the long-range imaging radar (LRIR) represents a category of radars known as inverse synthetic-aperture imaging radars. These radars use target motion instead of antenna motion to generate the synthetic antenna aperture. The LRIR operates in the X-band and has the range resolution that 1-GHz instantaneous bandwidth can provide (9.5-11.5 GHz). (The processed-signal equivalent of a transmitted pulse from LRIR has duration of 1 ns, which means it occupies approximately 1 ft of radar range extent.) And it has the cross-range resolution that the beam width of a 120-ft-diameter parabolic reflector can provide. Its transmitted power and receiver noise-floor characteristics are such that it can produce radar images of objects in geo-synchronous orbit 22,000 miles away.

The transmitter's RF high-power amplifier, which is shown in Fig. 11-19, comprises four VTX-5681 coupled-cavity TWTs. These tubes are rated at 100 kW of peak power and 30 kW of average power. (Originally, the tubes were supposed to be capable of 50 kW average power, but this requirement was downgraded to 30 kW as part of a cost-reduction modification.) Worthy of note is the high-power waveguide system, which was constructed in a special waveguide size, WR-102 (0.102-in. internal spacing between narrow walls instead of the 0.09-in. spacing for the standard WR-90 waveguide for this frequency band). WR-102 is the largest waveguide that can be used without incurring higher-order mode problems. The input to the high-efficiency, multi-mode tracking-antenna feed is in the form of eight equi-amplitude, equi-phase signals, which presumably could have been obtained by splitting the outputs of the four TWTs into two channels

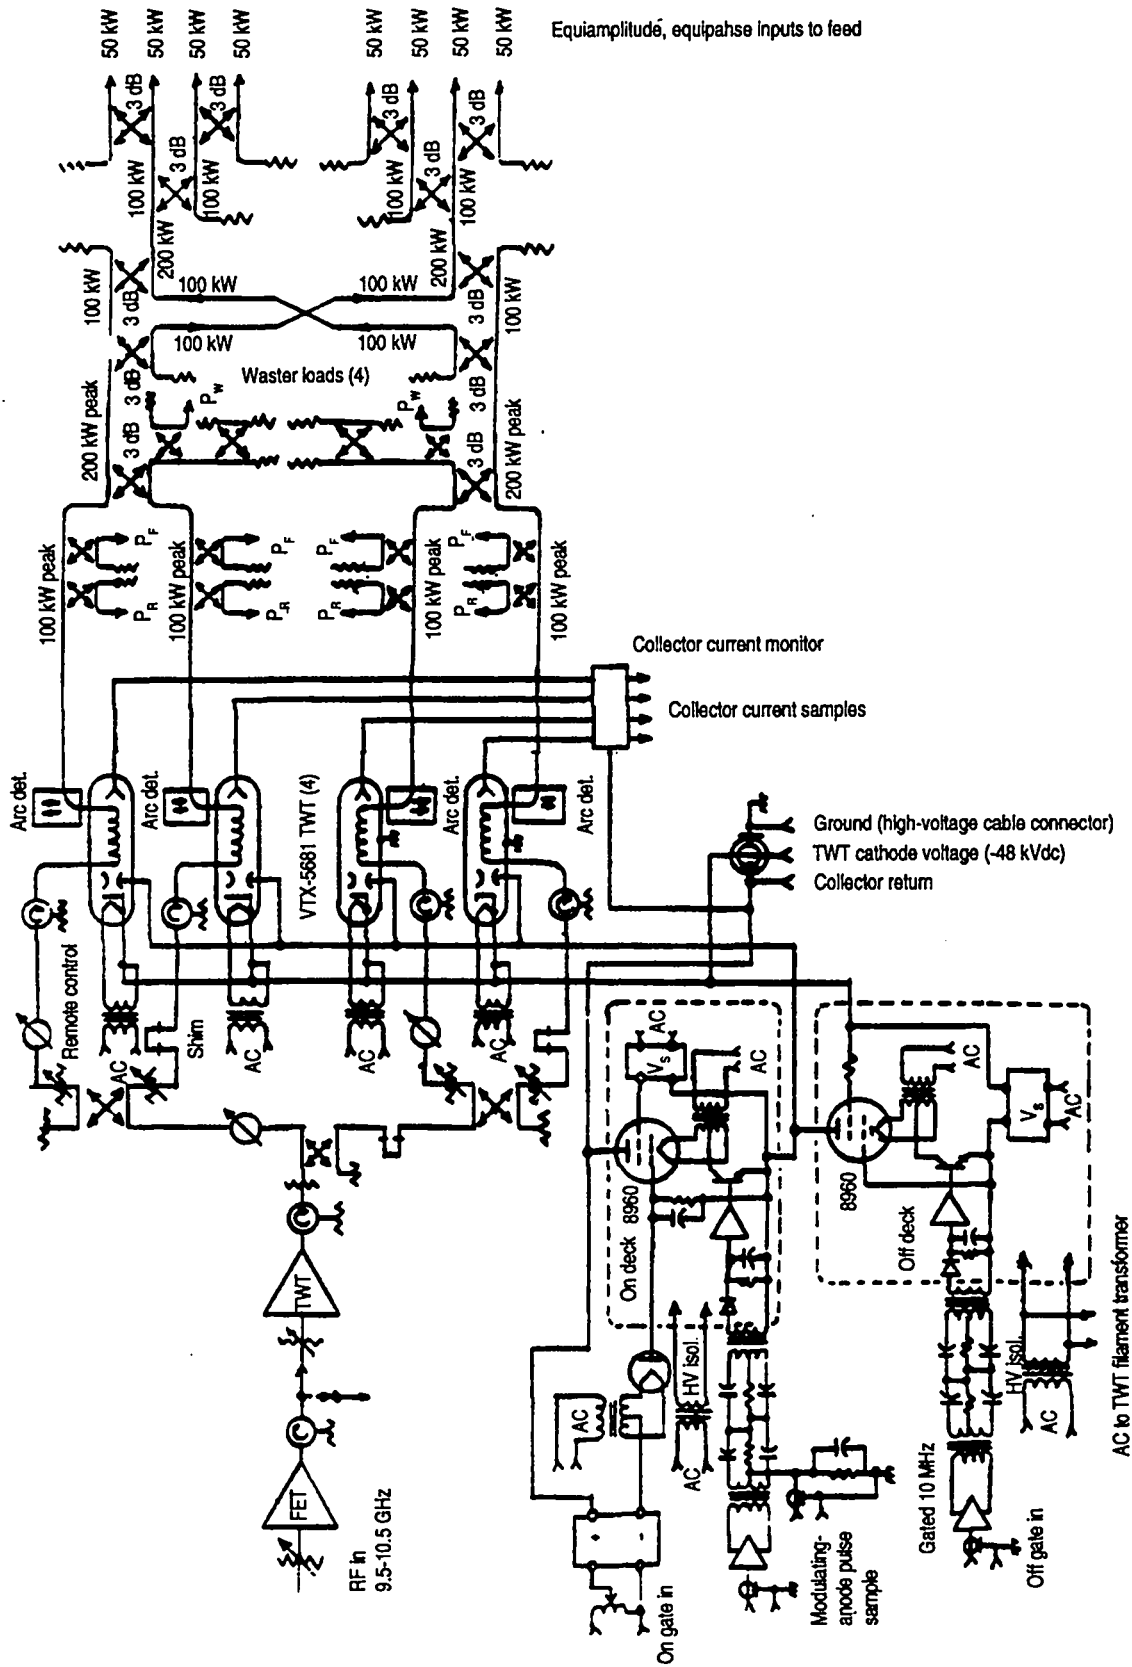


Figure 11-19. High-power RF amplifier and modulating-anode pulser of LRIR transmitter.

each. Doing so, however, would not have yielded eight channels of sufficient amplitude and phase identicalness over the entire 1-GHz bandwidth. (The differences in phase- and amplitude-transfer characteristics between individual TWTs are too great.)

To overcome this problem without combining the total output power in a single waveguide segment, five stages, or ranks, of combiners are used. The output of each pair of tubes is combined in the first rank of 3-dB hybrid combiners and then immediately split into two channels again in the second rank of 3-dB hybrids, which produces twice the power of a single TWT in each of two waveguide segments. These signals are then cross-combined, with one channel from each original pair of tubes combining with the remaining channel from the other pair in the third rank of hybrids, again producing twice the power of a single TWT in each of two waveguide segments. Now the composite signals in the two segments are identical to the extent that the passive waveguide components are identical. And at each frequency throughout the pass band, the composite signals are the average of the characteristics of the four TWTs. These signals are finally split into the requisite eight channels by the next two ranks of hybrids. Needless to say, all the high-power waveguides are water-cooled. The stainless-steel flanges compress soft-copper gaskets at each waveguide junction. The number of such junctions is minimized because the entire splitter/cross-combiner assembly is integrated.

The TWT electron guns, which operate at cathode voltage of just below 50 kVdc, have nearly full-voltage modulating-anode beam-current control to gate their 11-A peak-beam current. The modulating-anode pulser, as shown, is of the grid-catcher type. It has an active pull-up and pull-down and uses the type 8960 glass-envelope tetrode, which has adequate anode dissipation to permit repetition rates up to 2000 pps at the shortest pulse durations. The maximum pulse duration is 50 ms (50,000 μ s), which represents a high-stress condition for the TWTs but not for the modulator. The entire front end of the transmitter is air-insulated and, along with the receiver preamplifiers and downconverters, is mounted in an 8-ft x 8-ft x 12-ft enclosure that weighs 7000 lb. The feed horn, which is affixed to one end of the enclosure, is hoisted to the rear of the 120-ft parabolic reflector and slid forward so that the feed protrudes through a hole in the center of the main reflector and illuminates a subreflector that is part of the Cassegrain optics.

When originally put into service, the modulator was unique in that its low-level signal-coupling to the "on" and "off" decks was accomplished by means of capacitively coupled, balanced RF transmission lines. The inputs to them were in the form of pulse-modulated bursts of 11-MHz RF. The modulation envelopes corresponded to the desired modulator pulse-voltage output timing. (In effect, this was an update of the original BMEWS transmitter RF signal-coupling.) The balanced nature of the signal-coupling and the high ratio of the 11-MHz signal to the highest frequency component in the output pulse provided adequate common-mode signal rejection, even though the RF source was a computer-clock integrated circuit (IC). This IC, unlike the powerful 5-MHz transmitter used in the BMEWS-type transmitter, interfered with nothing else in the system. The capacitance of the signal link to the "on" deck and its voltage-equalizing shunt

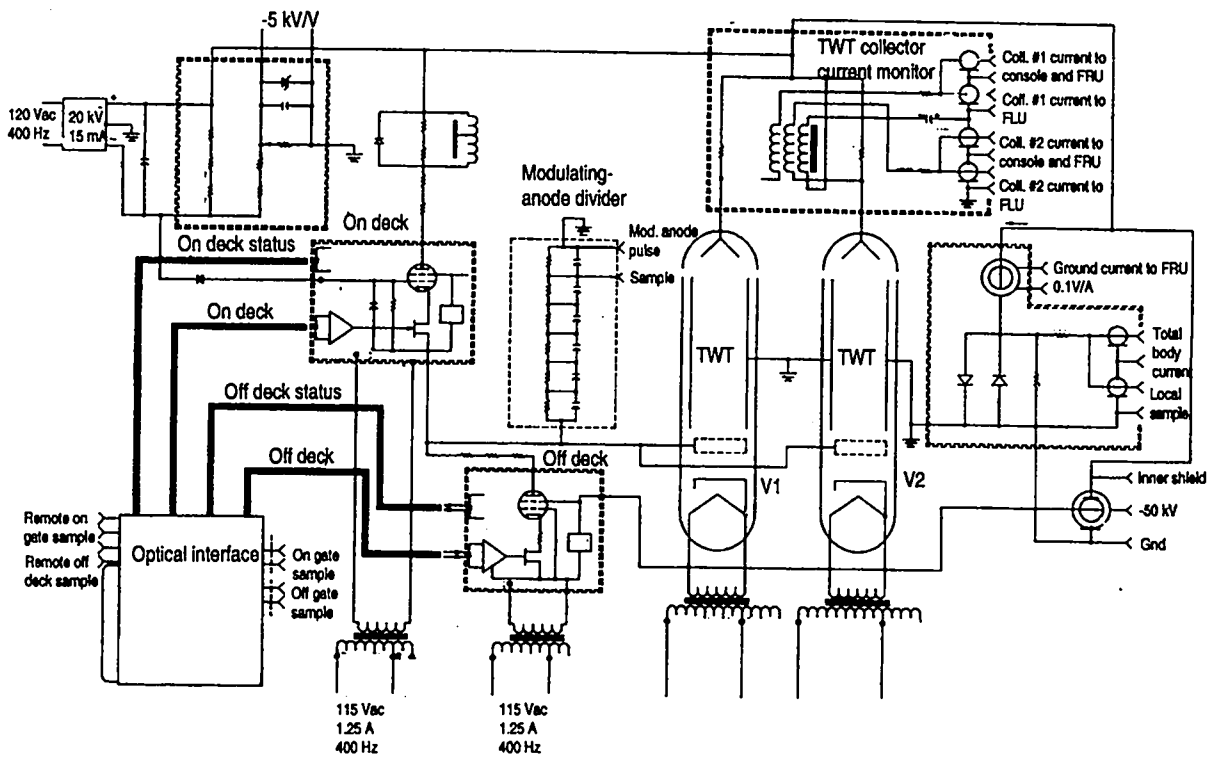


Figure 11-20. Simplified schematic diagram of modulating-anode pulser and dual-TWT high-power amplifiers of HAX transmitter.

resistance also doubled as part of the frequency-compensated voltage divider, which provided a direct-coupled sample of the modulator-output voltage. This original signal-coupling method has since been supplanted by modern fiber-optic signal links.

11.4.5 The modulating-anode pulser for the Haystack Auxiliary Radar (HAX)

Nearing completion in early 1993 is the Haystack Auxiliary Radar (HAX), which has its own dedicated, full-time 40-ft Cassegrain-optics antenna. (The LRIR described above must time-share its antenna with a number of box-mounted, physically interchangeable radio-telescope receiving systems.) HAX will have twice the bandwidth of the LRIR, 2 GHz, and operate in the K_U band. Its high-power RF amplifier consists of a pair of state-of-the-art coupled-cavity VTU-5692-type TWTs, which have re-entrant, double-staggered-ladder RF circuits. These TWTs are capable of as much as 60 kW peak-power output at 30% duty factor.

Along with the multi-channel low-noise receiver components, the entire front end of the transmitter, shown in Fig. 11-20, must share a fixed equipment compartment mounted behind the parabolic reflector. This compartment is quite a bit smaller than the RF box for the LRIR system. Despite the size limitation, however, the performance requirements for the modulating-anode pulser are very nearly the same as those for the LRIR transmitter modulator. The topology used is the same—active pull-up and pull-down, grid-catcher amplitude control—but everything has to be smaller.

The switch tube chosen is the highly capable 4PR250C. But at the maximum PRF of 2000 pps, its 250-W anode-dissipation rating is not enough. To help physically and electrically, the dimensions of the modulator decks are minimized, thus reducing the self-capacitance component of the modulator load capacitance and allowing the decks to fit in the available space. (This solution was facilitated by using 400-Hz primary power from a remotely located electronic static inverter.) Using a 400-Hz power supply allows the designers to use much smaller magnetic components, especially the filament transformers. Using smaller components permits deck dimensions of 6 in. on a side—just large enough for the tube socket. The other key enclosure-mounted components, including the low-capacitance power-isolation transformer for the “on” deck, the high-voltage isolation transformers for the “off” deck, and the two TWT filaments, can be reduced in size as well. (See Fig. 11-21.) To share the power dissipation with the switch tubes, current-limiting series resistance is used in the anode circuits of both switch tubes. The resistance in series with the “on” tube is great enough to noticeably affect the rise time of the output pulse. Half of the resistance, therefore, is shunted by an inductor that “shunt-peaks” the response, giving a rise time of less than 10 μ s.




Figure 11-21. Arrangement of components and subassemblies of HAX transmitter pulser.

12. Control-Grid Modulators

Of course, the size, weight, complexity, and expense of a pulse modulator can be minimized if the voltage and power required by the current-control electrode of the microwave tube can also be minimized. At the present point of evolution of the electron tube, the high- μ , high-transconductance grid is the control electrode that best achieves these goals. Regardless of the current intercepted by the control electrode, the switching losses, or the energy that must be dissipated in charging and discharging the capacitance associated with the electrode, for the control-grid tube, each pulse will vary as the square of the voltage swing. All major parameters being approximately equal, the capacitance associated with a control grid is about the same as that for a modulating anode. But the voltage swing for a grid may be only a small fraction of that required by a modulating anode, and the switching losses will be lower by the square of that ratio. Therefore, for high-repetition-rate operation—typically above 5000 pps—a control grid is almost mandatory, regardless of duty factor.

To be effective, however, the grid must be as close to the cathode as possible. But this presents problems. The cathode is hot and so it will radiantly heat the grid. The cathode also emits electrons, which are intercepted by the grid. Both of these factors will limit the average-power capability of the control-grid tube. But just as triodes and tetrodes have grid geometries that can minimize beam interception, so does the control-grid tube. The most popular geometry is the "shadow grid." The shadow grid is not the grid itself but is a masking electrode that is located near, on, or even embedded in the surface of the cathode. It inhibits emission from the cathode in those regions where electrons could directly strike the control grid, thus casting an electron "shadow" on the control grid.

Even if the shadow-grid concept were completely effective in keeping electrons from hitting the grid, and even if the grid itself were capable of dissipating unlimited power, there would still be an average-power penalty that the gridded-tube must pay. This is illustrated in Fig. 12-1, which contrasts computer-derived electron trajectories for an electron gun using a modulating anode with trajectories for a shadow-gridded gun. The electrostatic lens produced by the modulating-anode geometry is at least theoretically capable of generating a perfectly collimated, area-convergent electron beam; there are no undulations or "scallop-ing" tendencies of the beam as it enters the narrow beam tunnel and comes under the influence of the magnetic focusing field. (This is not to say, however, that the magnetic field cannot penetrate into the gun region. When it does, and if its strength is tapered so as to be congruent with the electrostatic lens, what results is called "confined-convergent flow," which results in the highest-quality electron beams achievable.) The gridded gun, however, produces an electrostatic aberration almost immediately; it introduces beam scalloping that no amount of magnetic focusing can completely mitigate.

Beam interception by the RF interaction circuitry of the microwave tube will always be higher for the gridded gun than for the diode or modulating-anode gun. This problem is often compounded by designs that use periodic perma-





Figure 12-1. Electron-beam formation in modulating-anode and gridded electron guns.

ment-magnet focusing, which also has imperfect polarity transitions along the beam. Added to this problem is the fact that the electrostatic lens is even less perfect for voltage transitions between grid and cathode, such as during the rise-and-fall times of the control-grid voltage pulse. This phenomenon is shown in Fig. 12-2 for a gun with a typical grid base. (The grid voltage for beam cutoff is -1% of beam voltage; grid voltage for full beam current is +1.8% of beam voltage, therefore the total grid voltage swing is 2.8% of beam voltage.) Note how the total cathode current splits between the collector and body as the grid voltage

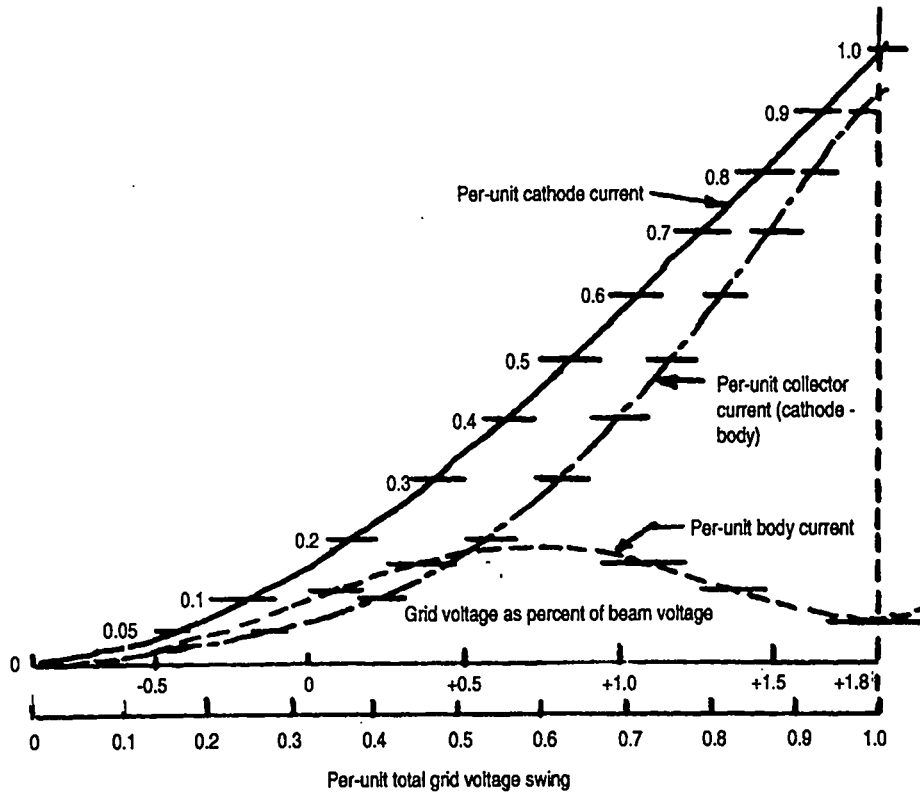


Figure 12-2. Typical transfer characteristic of shadow-gridded gun.

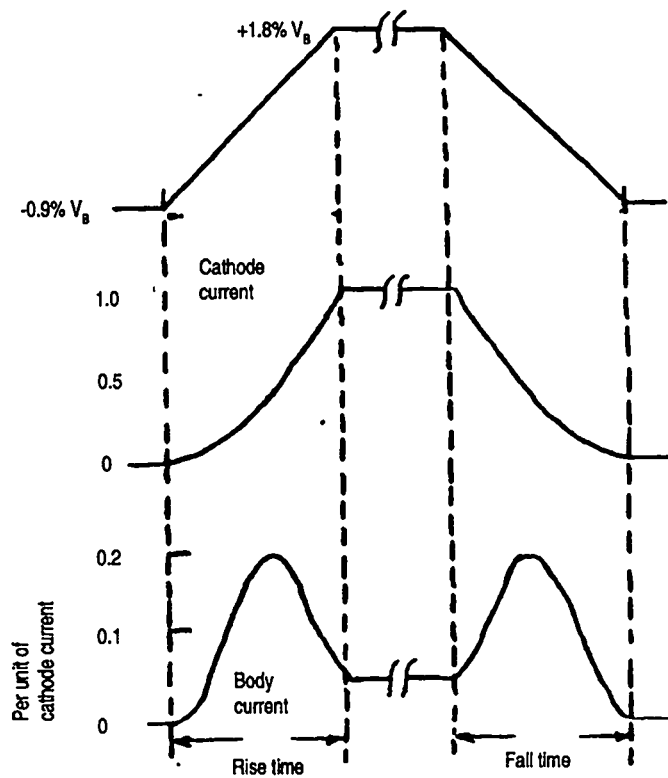


Figure 12-3. Time-domain voltage and current waveforms for tube with shadow-gridded electron gun.

swings from full-off to full-on. At full current, the portion intercepted by the body is only 5%. At about 60% of the grid-voltage swing, however, the body current rises to a value that is 20% of the eventual full-cathode current, which, at that point, is almost equal to the current that arrives at the collector. Beam transmission at that point is only slightly greater than 50%, whereas it will reach 95% at full current.

This is of no great consequence in long-pulse, low-repetition-rate service, where the percentage of the total "on" time devoted to the rise-and-fall portions of the pulse is small. But one of the reasons for using a control-grid geometry in the first place is to facilitate operation at a high repetition rate. Given the average-power constraints, the duty factor is usually only modest (typically less than 1%), which means that at high repetition rates—and rates up to 100,000 pps are not unheard of—pulses will tend to be short, and rise-and-fall intervals will be significant fractions of the flat-top durations. This situation gives rise to waveforms similar to those shown in Fig. 12-3. If we assume a linear rise and fall of grid voltage, then cathode current, varying as the $3/2$ power of instantaneous voltage, will approximately follow the shape shown. More important, however, is the shape of the intercepted body current in the form of "rabbit-ears." As the flat-top duration gets smaller, the ratio of the average value of body current to the average value of cathode current can be 2 to 3 times as great as it would be for the same useful (flat-top) duty factor, if the ratio was obtained with longer pulses at lower repetition rate. This problem, mitigated by faster grid-voltage rise-and-fall performance, theoretically disappears altogether with vertical-sided pulses. Vertical-sided pulses, however, cause infinite di/dt , subjecting any inductance in the cathode-current loop to infinite transient voltage, which has also been known to cause problems.

12.1 Grid-modulator topologies

Using the rule of thumb that grid-voltage base (or the required grid voltage swing from fully cut-off to fully turned-on) is less than 3% of the operating beam voltage, we can see that voltage swings of less than 1000 V will suffice for tubes operating at less than 30 kV of beam voltage, which includes the vast majority of tubes. (We will later discuss a notable exception.)

Today's transistor technology supports the design of single-device circuits that will switch 1-kV pulses—with not much safety factor, to be sure. This situation means that solid-state designs dominate the field of grid modulators. The availability of reliable and inexpensive fiber-optically coupled signal-links that are immune to noise and interference has made direct-coupled modulator designs far more practical than they were not so long ago. Nevertheless, many grid modulators use pulse-transformer signal coupling of the entire grid pulse, which is generated by a ground-referenced pulse generator. Some have even used capacitive coupling. (When duty factor is low, as it usually is, dc restoration of the coupled pulse is not required.) With capacitive coupling to the grid from a ground-referenced pulse generator, there will be a natural tendency for an intercepting-type grid to clamp the pulse voltage at zero bias because the grid-cathode circuit of the microwave tube will function as a forward-biased clamping diode. The average value of this rectified pulse current will develop a voltage

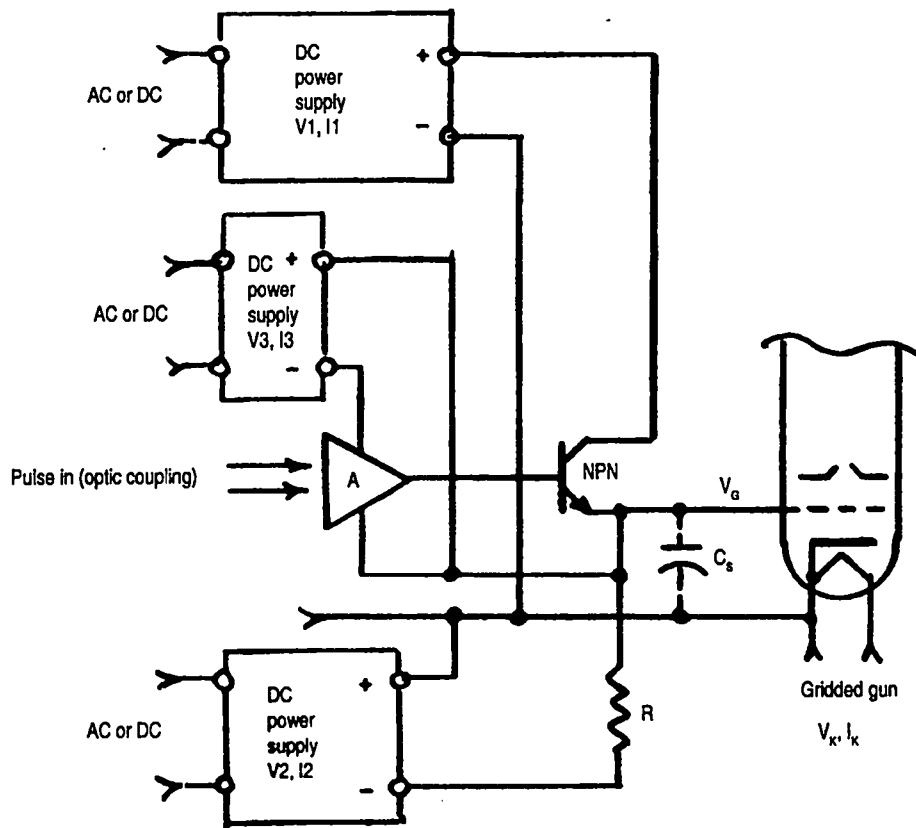


Figure 12-4. Grid modulator with bootstrap active pull-up, passive pull-down.

across whatever shunt resistance is connected in the circuit between grid and cathode. This voltage is called "grid-leak" bias. The bias voltage automatically adjusts itself so that the positive excursion of the capacitively coupled pulse just reaches zero grid-cathode voltage. In order to drive the grid positive, it is necessary to provide a low-impedance sink at the desired fixed bias to absorb the rectified grid current.

Common-mode rejection is not a great problem, because the microwave-tube cathode voltage, which is the return reference for the grid pulse, is nominally non-variant with respect to ground. (However, it will experience a step-voltage drop due to pulse-cathode current flowing through the fault-current-limiting resistance and a linear droop due to charge removal from the energy-storage capacitance.) Fast rise is made difficult by the leakage inductance of the pulse transformer, which can be physical large, too. The higher the microwave-tube beam voltage, the greater will be the insulation required between the primary and secondary of the pulse transformer—and the more difficult it becomes to minimize leakage inductance between windings. The coupling of long pulses depends upon the magnetizing inductance and core-saturation characteristics of a pulse transformer and/or the magnitude of the capacitance value of a coupling capacitor.

Transformer-coupled grid pulsers do have an advantage in that greater grid voltage can be obtained at the expense of proportionately greater pulse-generator output current by using the appropriate value of the secondary-to-primary turns

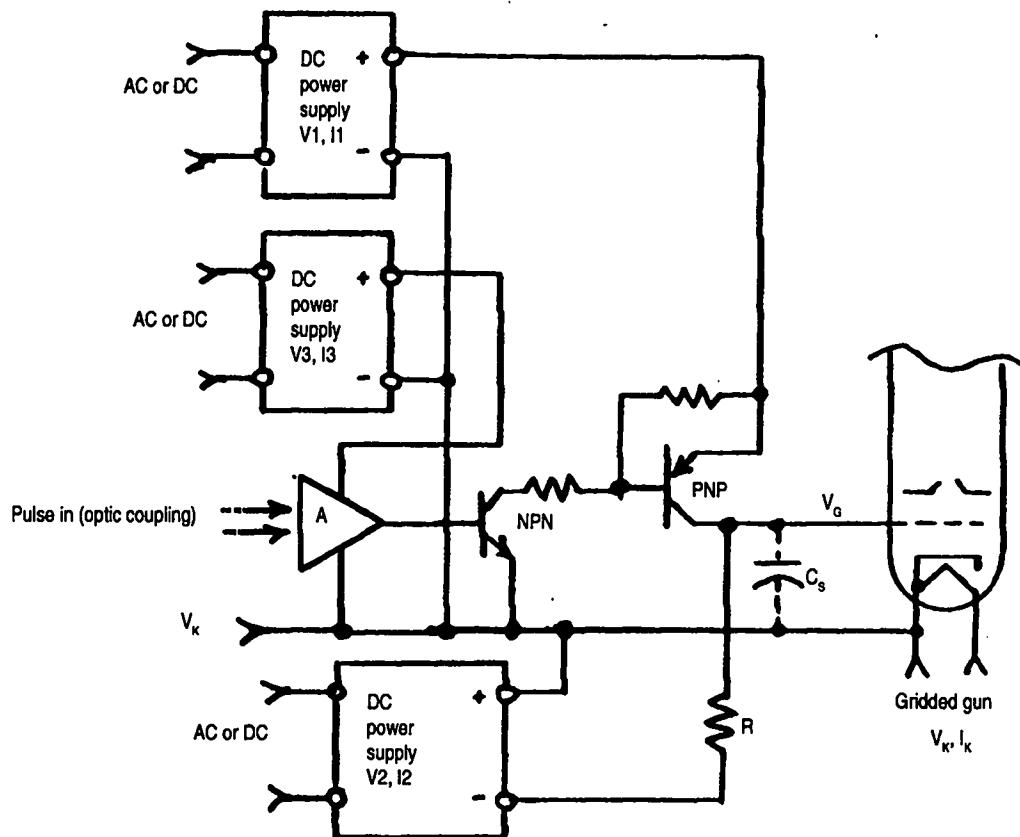


Figure 12-5. Grid modulator with open collector active pull-up, passive pull-down.

ratio. Fixed dc interpulse grid bias must also be generated by a power supply referenced to the microwave-tube cathode. The input to this supply can usually be obtained from the microwave-tube filament supply, which is already isolated from ground for the full dc cathode voltage.

Direct-coupled grid-modulator circuits do not suffer from duty-factor-induced level shifts or pulse-top duration limitations. The specific topologies of these modulators will now be described.

12.1.1 Active pull-up, passive pull-down

As before, we will begin our discussion with the simplest type of topology: the single active switch, as shown in Fig. 12-4. However, unlike the modulating-anode pulsers, passive pull-up is not a practical option for grid electrodes because of the finite intrapulse grid current that the grid pulser must supply. This current must flow through the pull-up resistor. The passive pull-down resistance, as before, determines the fall time-constant of the modulator pulse, RC_S , where C_S is the stray capacitance of the grid to cathode and focus electrode. The active pull-up switch is shown schematically as an NPN bipolar transistor, but is more likely to be a high-voltage N-channel MOSFET. The output is taken from the emitter, or source, but the transistor is not an emitter or source follower. It is a true bootstrap connection. The optic receiver and pre-driver, lumped together as amplifier *A*, and their housekeeping dc power supply (*V3, I3*), have common returns that are connected to the grid-pulse output. They should have low self-

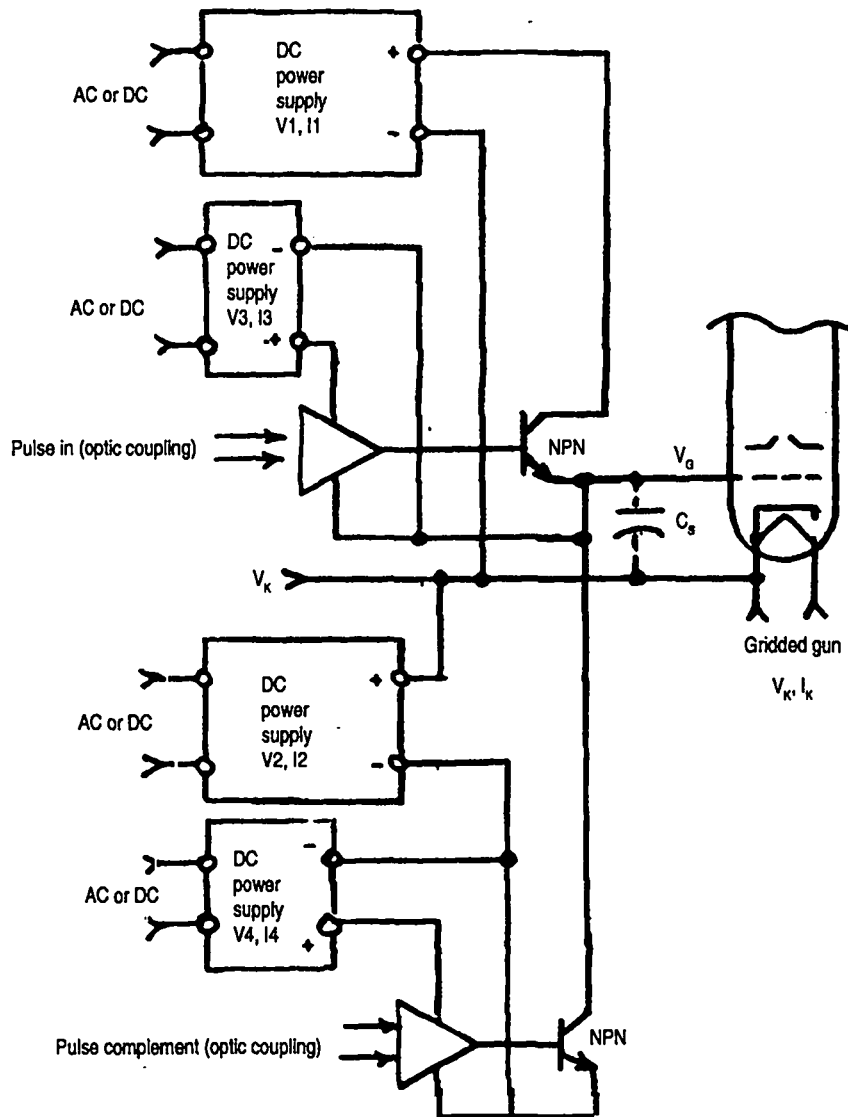


Figure 12-6. Solid-state grid modulator with active pull-up (bootstrap type) and active pull-down.

capacitances because they add to C_S in determining rise-and-fall times.

There are two other dc power supplies whose return leads are connected to the microwave-tube cathode. Supply V_1 must be equal to the sum of the voltage drop of the pull-up switch (usually negligibly small), which is approximately 1.8% of the beam voltage, V_K . And it is equal to the amount by which V_G is pulled positive. Supply V_2 is equal to the interpulse cut-off bias voltage, approximately 1% of V_K . The voltage across the transistor switch during the interpulse interval is $V_1 + V_2$, or approximately $0.028 V_K$, which is also the voltage across R during the pulse. The average currents, I_1 and I_2 , have two components: $0.028 V_K/R \times$ duty-factor, and $0.028 V_K \times C_S \times$ repetition rate. Assuming the conduction-voltage drop of the transistor switch is negligibly small, its average-power dissipation is $1/2 C_S \times (0.028 V_K)^2 \times$ repetition-rate. The average-power dissipation in the pull-down resistor is

$$\frac{(0.028V_R)^2}{R \times \text{duty factor}} + \frac{1}{2} C_S (0.028V_R)^2 \times \text{repetition rate}.$$

If the duty factor is small, as it often must be for a gridded-microwave tube, this simple topology is often the best.

An internal gun arc in the microwave tube can result in almost-instantaneous destruction of the transistor switch if it occurs between tube body and grid. Such an arc can also destroy the grid. So microwave tubes with gridded guns are designed to minimize the possibility of an internal arc path that terminates directly on the grid itself. The focus electrode, usually electrically connected to the cathode, is positioned within the gun and shaped so as to function as a lightning arrester, or arc shield, so that an internal arc will form between the body and focus electrode.

Figure 12-5 shows a grid pulser with the same basic topology, but it makes use of either a PNP bipolar or P-channel MOSFET output stage. This refinement gives the open-collector output connection a slightly more robust nature, thereby improving, at least marginally, the circuit's tolerance for microwave-tube internal arcs.

12.1.2 Active pull-up and pull-down grid pulser

When duty factor becomes too great or pulse fall time becomes very short, or both, pull-up and pull-down must both be active. The power supplies for $V1$ and $V2$ have similar roles as before. Only $I1$ is duty-factor dependent, having two components: (peak-pulse grid current) \times duty factor; and $(V1 + V2) \times C_S \times$ repetition rate. Current $I2$ is $(V1 + V2) \times C_S \times$ repetition rate. Average-power dissipation in both transistor switches, assuming small voltage drop, is $1/2 C_S (V1 + V2)^2 \times$ repetition rate, and average microwave-tube grid dissipation is $V1 \times$ (peak-pulse grid current) \times duty-factor.

As shown in Fig. 12-6, an additional housekeeping power supply, $V4$ and $I4$, is required for the optic receiver and pre-driver for the pull-down switch, but self-capacitance is of no concern because it is not a bootstrap circuit and there is no common-mode pulse voltage.

12.1.3 Single-switch active pull-up and pull-down grid pulser

How can you have active pull-up and pull-down with only a single electronic switch and no cutoff bias power supply? The answer is shown in Fig. 12-7. It is yet another variant of the cascode-connected cathode drive that has already shown up here and there. The active pull-up is literally active pull-down because, in order to initiate pulse-current flow, it is the cathode of the microwave tube that will be pulled down by driving the NPN transistor (or N-channel MOSFET) into conduction. The grid of the microwave tube is connected to the positive terminal of a dc source whose output voltage is equal to the positive grid voltage required for nominal peak beam current ($V_G = 0.018 V_{BM}$). During the interpulse interval, the voltage at the junction of the transistor collector (or drain) and microwave-tube cathode automatically rises to a value equal to the cathode-grid voltage required for beam cutoff (1% of V_{BM}) plus the grid power-supply voltage (1.8%

V_{BM}). This amounts to the same total of 2.8% V_{BM} that was switched by the previous grid modulators. When the transistor switch is turned on, current through it discharges the stray capacitance C_S , pulling the tube cathode voltage toward the V_{BM} bus. The rate at which this happens is determined by the total available switch current or conduction-state resistance. Note that the capacitance that must be charged and discharged is not the same as when the grid is pulsed. This includes the capacitance of the filament supply V_F , which should be designed for minimum value.

During the flat-top intrapulse interval, the transistor switch must conduct the entire beam current of the microwave tube. However, it is this current *after* the switch is turned off that provides the "active" pull-down, which is literally pull-up, by recharging the stray capacitance to its initial interpulse value. As discussed earlier, neither this current nor an equivalent series resistance is constant during the recharge, or trailing-edge, portion of the pulse, so that the waveform will be neither linear nor exponential.

The fault tolerance of this circuit connection depends upon the transient overvoltage capability of the transistor switch, which must have external transient suppression, shown as an MOV. If the microwave-tube electron gun has its grid shielded from internal arcing, the arc will terminate on a cathode-potential electrode, and the arc-current outlet is through either the transistor switch or the transient suppressor. If the circuit is to survive, the arc current had better flow through the transient suppressor, whose clamp voltage, even for short-duration residual overvoltage due to series-inductance effects, must be coordinated with the transistor safe-operating-area restrictions.

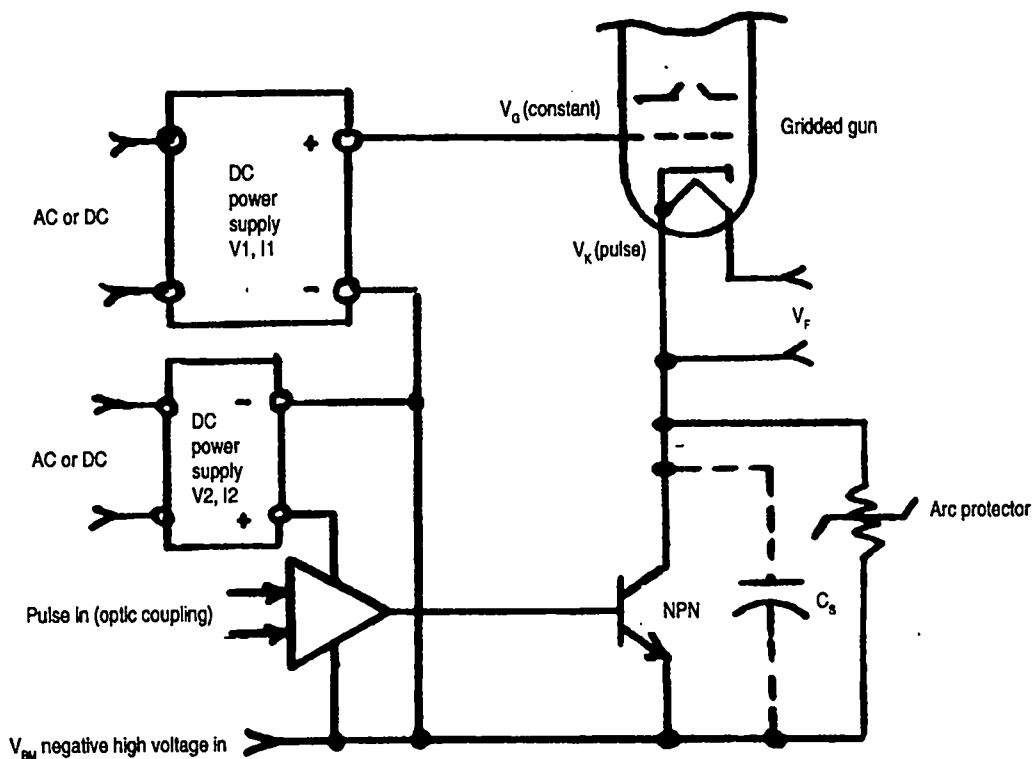


Figure 12-7. Another solid-state grid modulator with active pull-up and pull-down, but only one control switch.

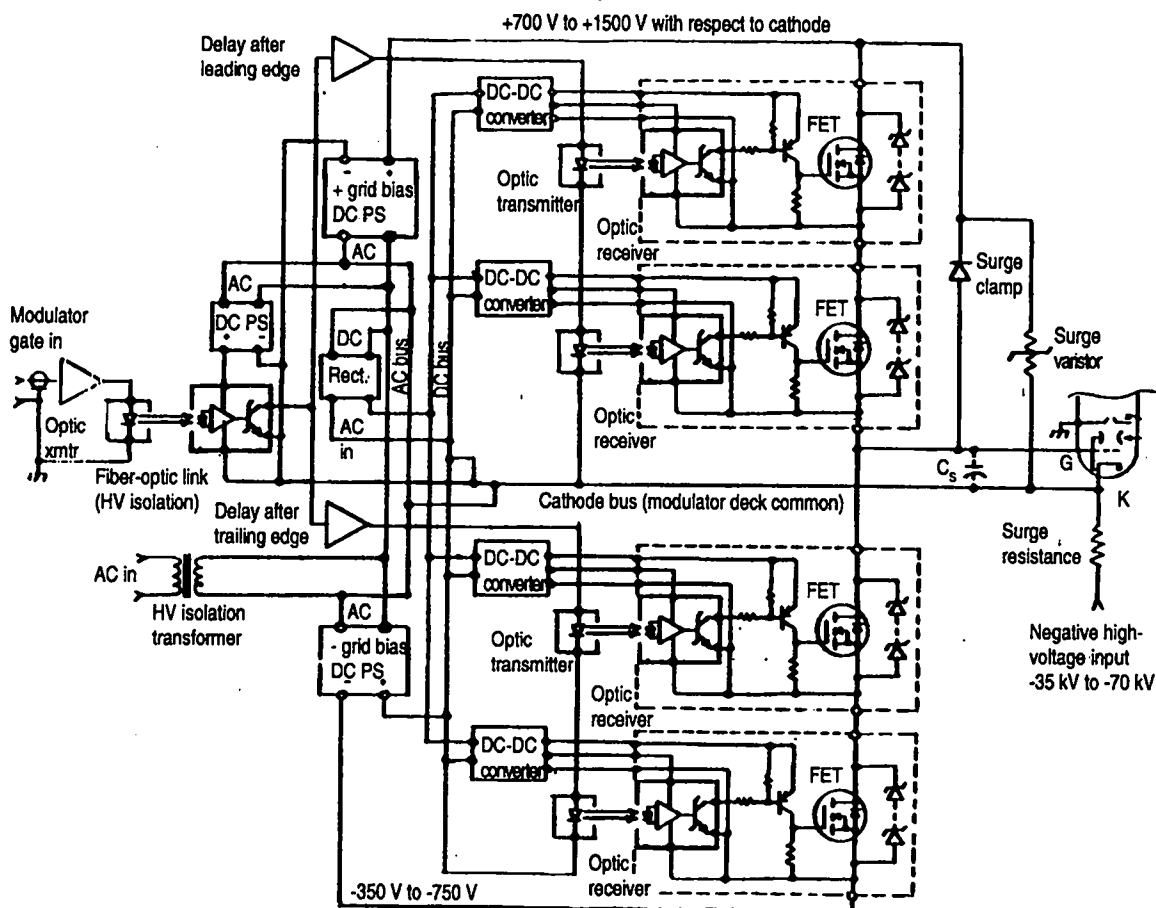


Figure 12-8. Cascaded solid-state modulator with active pull-up and pull-down for gridded guns operating at higher grid bases.

12.1.4 Cascaded solid-state active pull-up and pull down

Not all gridded-gun microwave tubes operate below 30-kV beam voltage. Those that don't may have grid bases that exceed the safe-operating conditions of single transistor switches. There is nothing, except circuit complexity and expense, that prevents the use of series-cascading transistor-switch modules to accommodate the switching voltage requirement of virtually any gridded electron gun. An example of such a stacked-module topology is shown in Fig. 12-8, which is but one of many practical implementations.

The circuit shown is expanded to only two modules each for the pull-up and pull-down segments of the pulser. A single optically coupled signal link is used to inject a low-level modulator gate to an electronics deck whose common plane is connected to the cathode of the microwave tube. In the absence of an optical signal, the quiescent, or interpulse, conditions are obtained. During this period, the current is driven through the series-connected optical transmitters associated with the pull-down modules, thus maintaining the FET switches in the conducting state. (Note that this is not a "fail-safe" implementation, because active pull-down throughout the entire interpulse, or stand-by, interval requires the presence of current in the optic-transmitter loop, fiber-optic continuity, and fiber-optic

connection to each pull-down module optic receiver. It would be preferable if all modules in the complete pulser were identical because using modules with complementary logic would not be "fail-safe" during the active pull-up. However, pull-down resistance connected between microwave-tube grid and the negative-bias source would mitigate the situation, especially for the stand-by condition.)

When the leading edge of the modulator gate is received at the pulser gate, current in the optic-transmitter loop of the pull-down portion is terminated without intentional delay. The start of current in the optic-transmitter loop for the pull-up modules is intentionally delayed by a few microseconds to assure the recovery of the FET pull-down switches to their non-conducting states. This is done to preclude simultaneous conduction in pull-up and pull-down paths, or shoot-through.

At the end of the modulator gate, the process is reversed. Current in the pull-up optic-transmitter loop is terminated with no delay, while intentional delay is inserted in the start-up of current in the pull-down optic-transmitter loop. The circuit then toggles back and forth between full-conduction and full-cutoff grid conditions, using as many series-connected switch modules as safety-factor considerations dictate for grid base and switch-voltage. The classic problem of series-connected electronic switches is the lack of simultaneity in turn-on and turn-off, which results in unequal voltage-sharing during the leading-edge and trailing-edge transient conditions. This is why each transistor switch is shunted by a Zener-diode string whose transition voltage is coordinated with the voltage hold-off capability of its associated switch. The reluctance of a switch to either turn on or turn off manifests itself in a stair-step discontinuity in the rise or fall wave shape. It does not result in destructive transient overvoltage. Each switch module requires an isolated, low-capacitance source of housekeeping voltage and current located between isolated input and output terminals. In Fig. 12-8 this is shown as a dc-dc converter. (Later, a type of converter that is nearly ideal for such service will be discussed.)

Fault tolerance for this type of pulser is greater than that for the single-switch pulser only because this modulator is used in situations where more grid-voltage swing is required, so it is inherently a higher-voltage device. On a per-unit voltage basis, however, its susceptibility is the same (higher grid voltage, higher beam voltage). Just as in the single-switch topology discussed above, a MOV surge-voltage clamp is used as the primary path for fault current resulting from a body-grid internal arc. The MOVs can be effective, but they add considerable load capacitance to the circuit. Where blazing speed is required, their loading effect must be considered at all points of the design. (They certainly cannot be tacked on as an afterthought without measurable performance penalty.) Spark gaps, which have much lower shunt capacitance, are often considered for fault protection. They have a shortcoming, too. Their internal plasma development has inertia, which means that the gap's sparkover voltage increases as the rate-of-rise of voltage increases. In any case, even the lowest-inertia spark gaps have a ratio of static voltage hold-off to transient-voltage breakdown of at least a 1.5:1. This means that for guaranteed protection the pulser must have at least a 1.5:1 voltage safety factor.

An alternative to direct shunt voltage-limiting is the surge diode connected between tube grid and the positive grid-voltage source, as shown. The only capacitance added is the reverse-biased diode capacitance. An internal arc to the tube grid, which will tend to pull the grid toward ground (forcefully), will forward-bias the diode as soon as the grid becomes more positive than the positive-source power-supply terminal, forcing fault current through the power supply (and reverse-biasing its internal rectifier diodes). If, however, the storage capacitance that shunts this supply is large enough and the total charge transfer from the beam power supply is not too large, the fault will be absorbed without damage to anything. The power-supply terminal voltage will rise by an amount $\Delta V = \Delta Q/C$, where ΔQ is the charge transfer from the beam supply energy-storage capacitor and C is the capacitance shunting the pulser positive-bias power supply. The peak current through the surge diode will be limited by the surge resistance in series with the microwave-tube cathode. The surge diode must be rated for at least this much peak current. In addition, the diode must be rated to handle the action integral

$$\int i^2 dt.$$

This integral, having the dimension of ampere²-second, or joules/ohm (which are the same), can be evaluated either by doing the time integral of the exponentially decaying fault current or by finding the total energy stored in the beam-supply system and dividing by the total series surge resistance.

12.1.5 Hybrid grid pulser

Even though today's "technologically correct" design strategy is to use solid-state components wherever possible—and the grid pulser is the only high-level modulator where their exclusive use is even remotely practicable—there are still applications where a combination of solid-state and vacuum-tube devices can yield an optimum design solution. For instance, there is nothing simple or inherently inexpensive about the cascaded-module pulser just discussed—especially as the module count becomes large. There are small, planar-geometry triodes such as the Y-540 that can hold off kilovolts (6 kV for the Y-540) and deliver short-pulse currents that are more than adequate for microwave-tube grid interception. Therefore, they could supplant a number of transistor switch modules.

The circuit of Fig. 12-9, an obvious adaptation of a previous modulating-anode pulser topology, shows such a hybrid. The pull-up, or "on," portion of the circuit uses a triode vacuum-tube high-voltage switch in a cascode connection with two MOSFET switches, which are optimized for separate functions. The upper FET, designated "switch," turns the triode on and off in response to the optically coupled modulator-gate timing signal. The lower FET, designated "clamp," is connected in a gate-catcher circuit in which its gate electrode is diode-coupled to an adjustable dc power supply that permits grid-pulse amplitude control from a ground-referenced operator adjustment. The anode supply for the pull-up triode has an output voltage that is approximately 1 kV greater than the maximum output-pulse amplitude, up to the anode-voltage capability of the triode. The low-current amplitude-control supply has a minimum output volt-

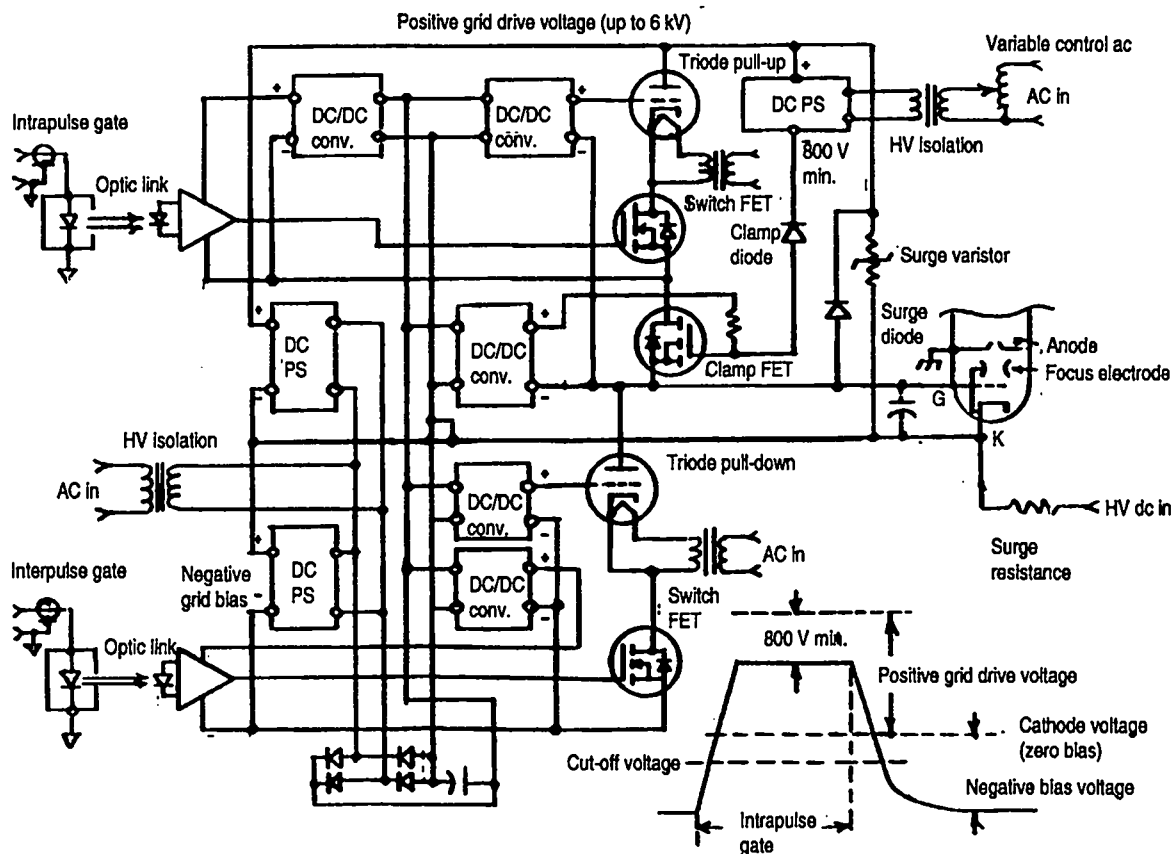


Figure 12-9. Hybrid solid-state/vacuum-tube grid modulator for high-voltage guns.

age of 1 kV because it bucks the anode supply. (The greater the amplitude-control supply voltage, the smaller the output-pulse amplitude.)

As in the comparable modulating-anode pulser topology, the timing-switch FET and the clamp-switch FET are shown as separate, isolated entities that are connected in series as a logical AND gate. With some compromise in overall performance (particularly affecting rise-and-fall times and reference-supply average current), the "clamp" FET can be also used to perform the timing function by connecting a small-signal NPN transistor switch between the FET gate and source. The NPN transistor conducts during the interpulse interval, removing FET gate drive, and is driven off for the duration of the modulator-gate signal, allowing the FET gate to rise until the clamp circuit pulls it back down again when the desired pulse amplitude has been reached (or slightly exceeded). Although this design results in a simpler topology, it is not "fail-safe," because the low-level transistor must conduct in order to prevent attempted pull-up. If it fails in the short-circuit mode, nothing potentially hazardous will occur. If, on the other hand, it is simply removed from the circuit, continuous pull-up current will flow. Unless the situation is sensed and corrected by status-monitor circuitry, it could damage the microwave tube, the modulator, or both.

12.1.6 "Inductive-kick" single-switch grid pulser

One answer to the question "what is the voltage across an inductor?" is "whatever voltage is required to maintain constant current flow through it." The operation of the high-voltage grid pulser shown in Fig. 12-10 is based on that answer. The figure's single electronic switch, which is made up of a number of series-connected transistors, conducts throughout the interpulse interval. In so doing, it connects the grids of the multiple microwave tubes (eight, in its most significant actual applications) to a -900-V grid-bias power supply. At the same time it conducts current from a 600-V, 300-mA constant-current power supply through a 16-H inductor.

When the trigger-amplifier/light-receiver receives an optically coupled signal corresponding to the beginning of a pulse, the electronic-switch transistors are turned off, opening the path for the 300-mA inductor current. With no place for the current to go (yet), the voltage rapidly builds up across the inductor, such that the voltage at the junction of the inductor and the electronic switch increases in the positive-going direction. This voltage becomes the leading edge of an output-voltage pulse that keeps rising until it exceeds the +2-kV output voltage of the clamp power supply, the conduction-voltage drop of the seven series-connected Zener diodes, and the bias-blocking diode in series with the clamp

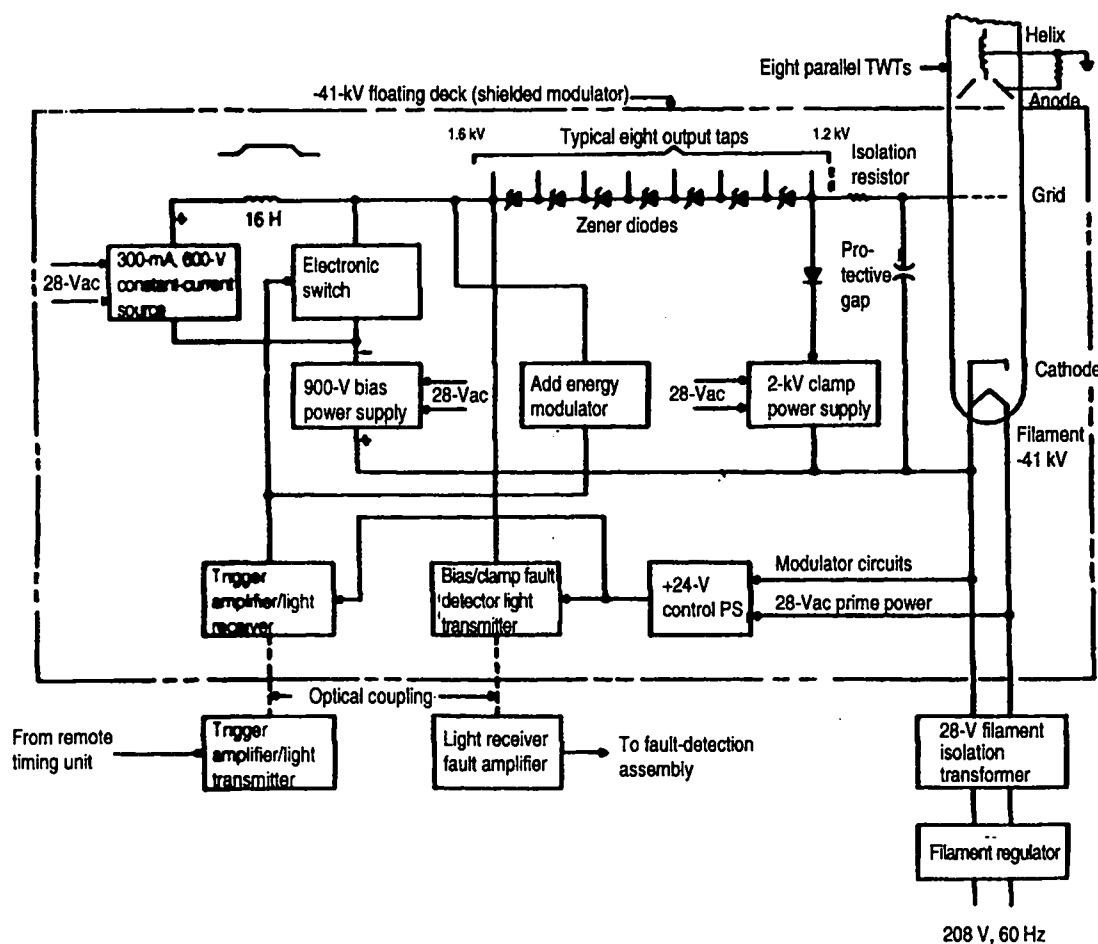


Figure 12-10. Inductive ring-up grid modulator for high-voltage gridded gun.




Figure 12-11. The 24-channel, single-output transmitter of the ALTAIR UHF radar.

supply. With a path thus re-established for the inductor current, the voltage no longer rises, and a flat-top pulse continues until either the volt-time product of the inductor core is exceeded or the electronic switch is turned back on, terminating the pulse and restoring the original negative-bias conditions at the pulser output. The inductor is large enough and has sufficient internal distributed capacitance so that the rate-of-rise of voltage is not quite as fast as desired. To increase the speed, additional current in the form of a sharp pulse is coupled to the output node from the add-energy modulator.

To satisfy any differing grid-drive requirements of the electron guns connected to it, the pulser has eight separate output voltages at 50-V intervals over a 400-V range. These can be tapped at the junction points between Zener-diodes. Note that the filaments of the electron guns operate from 28 Vac, coupled to the -41-kVdc floating deck through a filament isolation transformer, which also provides the prime-power input to the dc power supplies used by the pulser.

12.2 Transmitter applications of grid pulsers

As mentioned earlier, gridded-gun microwave tubes—primarily the helix-type TWTs—whose peak power is usually less than 5 kW are almost universally used in systems demanding a high degree of waveform versatility. One use for such a transmitter is as a radar jammer for electronic-countermeasures applications. The pulse trains produced by this transmitter may have pulse recurrence intervals of only a few microseconds and repetition rates up to 100,000 pps, while their flat-top pulse durations may be on the order of fractions of a microsecond. Pulse Doppler radar systems use the frequency domain in much the same way that long-range search radars use the time domain. To achieve low Doppler ambiguity, a transmitter must have a high repetition rate (for greater frequency deviation between PRF lines in the frequency domain), whereas to achieve low range ambiguity, a transmitter may have more time between successive pulses and a low repetition rate.

There are literally hundreds of designs of small, solid-state, high-speed grid modulators for the host of compact, light-weight transmitters used in military aircraft. On the other hand, there are only a handful of truly high-power transmitters that use RF power amplifiers with gridded electron guns. And in this group are some of the most powerful transmitters in the world.

The last class of grid pulser to be discussed is the one used in the most impressive of these transmitters: the ALTAIR (ARPA Long-Range Tracking and Instrumentation Radar). The ALTAIR UHF transmitter uses 24 TWTs rated at 200-kW peak power, 6% duty factor, and 2-ms pulse duration. The tubes are organized into three groups of eight tubes each. The transmitter is of interest not only because it uses three of the grid modulators described above, but because it is a multi-channel (24) transmitter that starts with a common RF drive signal in one RF transmission line and ends up with the outputs of all of the TWTs combined in a single-output waveguide. (This waveguide, in turn, feeds a 150-ft-diameter parabolic-reflector antenna system through a multi-mode tracking feed and Cassegrain optics.) Why is the number 24 of particular interest? Because it is not part of a binary progression, like 2, 4, 8, 16, and 32.

Because the transmitter's microwave-tube organization does not follow a bi-

nary pattern, it follows that power splitting and power combination cannot be accomplished by relying exclusively on networks comprising 3-dB couplers. In fact, as shown in Fig. 12-11, which depicts the simplified RF circuit of the 24-channel transmitter, the very first power-splitter hybrid has a coupling factor of 4.77 dB. This value was chosen because it produces a 2:1 power split, with 1/3 of the input power, P_{in} , going to the quadrature port, and 2/3 P_{in} to the in-phase port. After that, the remaining breakdown of the input signal into 24 equi-amplitude channels is accomplished with 3-dB hybrid junctions. The same hybrid junctions recombine high-power outputs of each tube of a group into one channel that is the sum of the group. A second similar junction sums the output of the other two tube groups. The last step of combining is accomplished in another 4.77-dB hybrid junction implemented in WR-2100 waveguide. The progression of 90°-phase-shift increments produced by the quadrature hybrids leads to the final in-phase condition of all of the voltage vectors at the output. This progression is listed in the figure as are the voltage amplitudes throughout the combiner tree. The final power output of the system can be determined by the output voltage, V_0 , as expressed in the following equation:

$$V_0 = 4V\sqrt{\frac{2}{3}} + 2\sqrt{2}V\sqrt{\frac{1}{3}} = 6\sqrt{\frac{2}{3}}V.$$

Power output, P_0 , can be defined as

$$P_0 = kV_0^2 = kV^2 \times 36 \times \frac{2}{3} = 24kV^2,$$

where kV^2 is the TWT output power of 200 kW. Final output power is 24 times 200 kW, or 4.8 MW.

Just as there can be three-cylinder engines, there can be non-binary combiners. Yet, how do we determine the required coupling value for the needed hybrid junctions after all of the binary-progression numbers have been factored out? The quadrature-hybrid criterion for zero waster-load power, shown in Fig. 12-12, can supply the answer.

Assume that two arbitrary input voltages, V_1 and V_2 , already have the requisite 90° phase differential at the two input ports of a quadrature hybrid, which has a voltage coupling factor of A such that the voltage coupled from an input port to the quadrature output port is 1 divided by A , as shown. If all of the power represented by the two input voltage vectors is to combine at the desired output port, then none of it can be coupled to the opposite, or waster-load port. The two voltage components incident on the waster load are already in phase opposition. If they are of equal amplitude, they will cancel out. To be equal,

$$V_1\sqrt{1 - \frac{1}{A^2}} = \frac{V_2}{A},$$

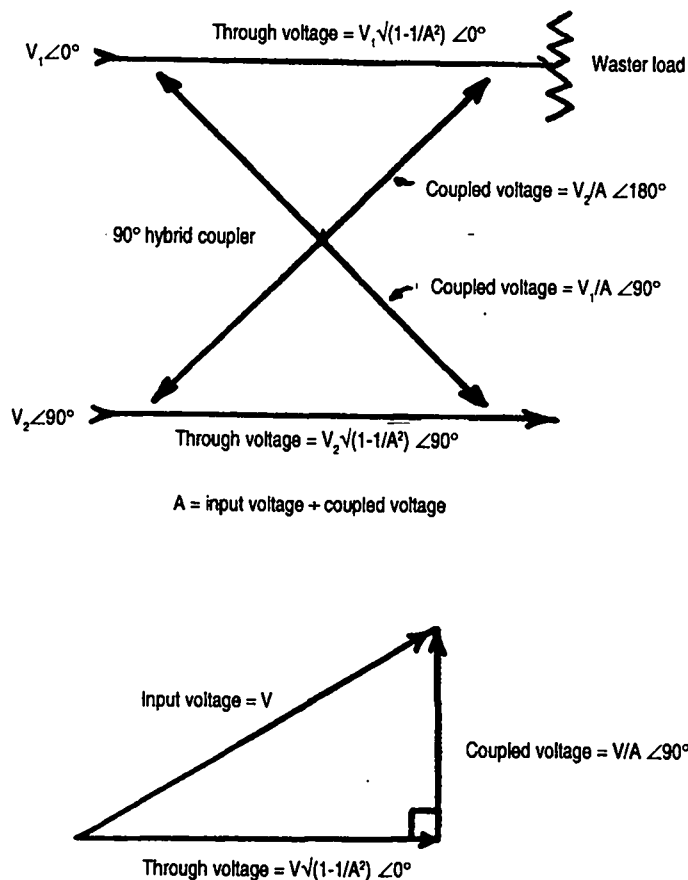


Figure 12-12. General relationships in quadrature combiner-coupler.

which gives

$$A = \sqrt{\frac{V_2^2}{V_1^2} + 1} .$$

V_2^2 and V_1^2 are proportional to the corresponding power levels P_2 and P_1 , so A can be expressed as

$$\sqrt{\frac{P_2}{P_1} + 1} .$$

Because A is a voltage ratio (presumably in a constant-impedance system), the coupling factor can also be expressed as

$$20 \log \sqrt{\frac{P_2}{P_1} + 1} \text{ dB} ,$$

or

$$10 \log \left(\frac{P_2}{P_1} + 1 \right) \text{ dB} .$$

For the case where P_2 is twice P_1 , the coupling factor will be

$$10 \log(2 + 1) = 4.77 \text{ dB} .$$

For the more common equal-power case, $P_1 = P_2$, the coupling factor is

$$10 \log(1 + 1) = 3 \text{ dB} .$$

13. Charge Diverters: the Electronic Crowbar

With the exception of most line-type pulser, all of the modulator types discussed require some high-speed means ($\sim\mu\text{s}$) of providing an alternative path for fault current that results from an internal arc in the electron gun of our microwave tube.

13.1 What happens when an electron gun arcs?

A typical circuit arrangement and fault scenario is shown in Fig. 13-1. It involves a microwave tube with a gridded electron gun. The presence of the grid, which comprises small-diameter wires, coupled with the possibility—even if only rare—that an arc can terminate on it, brings into play both components of possible arc damage. The first component is the energy dissipated in the arc itself, which is the time integral of the arc voltage drop and the arc current. The voltage drop of an arc between copper electrodes in vacuum, shown in Fig. 13-2, is almost independent of arc current, up to at least 1000 A. And it is small, less than 20 V. Beyond 1000 A, some experiments have shown the rising characteristic illustrated in Fig. 13-2, where arc voltage drop increases to about 50 V at a current of 10,000 A. It increases rapidly beyond that as the arc becomes unstable. Other experiments have shown low arc voltage drop for currents well beyond the apparent upper limit of 10,000 A. In any case, 20-V arc drop for up to 1000-A arc current can usually be depended on. With the arc voltage constant, the other factor is the time integral of the current,

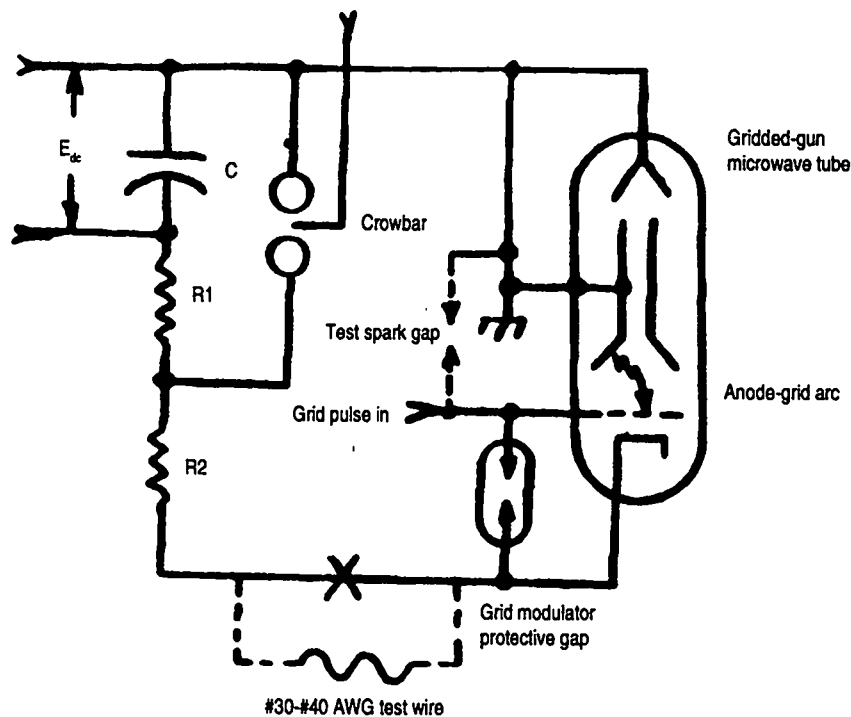


Figure 13-1. Typical fault scenario for gridded-electron gun.




Figure 13-2. Properties of an arc in a vacuum between copper electrodes.

$$\int i \cdot dt,$$

which is nothing more than the total charge transport through the arc.

Referring back to Fig. 13-1, the circuit elements that determine what the fault current will be are the surge resistors $R1$ and $R2$. Their total value must be V_{dc}/I_{FAULT} , where V_{dc} is the voltage across the storage capacitor C . If no charge-diverter path or crowbar is provided, then the charge transport through the arc will be the charge Q that is initially stored in C , where $Q = V_{dc} \times C$. The arc energy will be the product of the arc drop, 20 V, and charge Q . A commonly quoted permissible arc energy for microwave-tube electron guns is 50 J, although this value has been by no means rigorously derived. Nevertheless, using this criterion we could stipulate that in cases where stored charge is less than 2.5 C, no crowbar is required, and when it exceeds 2.5 C, one is—providing that the fault current is limited to 1000 A or thereabouts. Note that no mention has been made so far of the energy stored in the capacitor bank C . This is because all but a tiny fraction of the stored energy will be dissipated in either $R1$, if the crowbar fires in a timely fashion, or in $R1 + R2$, if it does not. (If there is no crowbar at all, $R1$ and $R2$ can become the same resistor.)

An electronic crowbar is often referred to as an energy diverter, which is certainly what we want it to do. However, it is no more capable of dissipating energy than the initial arc. The energy must be dissipated in the resistors. The crowbar switch becomes an alternative path for discharge current, diverting the charge from the initial arc. The current through the crowbar is limited by $R1$, which must be large enough that the diverted current does not exceed the capability of the crowbar arc. In some cases, $R1$ can be made large with respect to $R2$, in which case crowbar and fault currents are nearly the same.

A crowbar charge diverter of low-enough impedance and rapid-enough firing

time can limit the energy deposited in the initial gun arc to a tiny fraction of what it would have been. Often, however, this residual energy is actually too small to do any good in "processing" the voltage-hold-off capability of the faulting gun. Believe it or not, there are often good reasons for repetitive gun arcing—to control the growth of micro-projections, for instance. Micro-projections are metallic filaments that protrude from an otherwise-smooth copper surface. The tips of these projections have greatly enhanced electric-field intensities. An arc can burn these projections away. But if there is not enough energy in the arcs that the micro-projections can induce, then they will never be burned away. There is an optimum value of "let-through" energy for each individual case. Unfortunately, there is no analytical criteria for what that value might be. A high-speed electronic crowbar can always be delayed in its response, and some have been, with marked improvement in subsequent gun arc rate.

Other than the "50-joule" criterion, prior analytical determination of optimum crowbar performance has never been quite possible. But the experimental calibration of crowbar performance is. This procedure is often referred to as the "tin-foil" test, although 1-mil household aluminum foil is the experimental medium most often used. The test is performed using a calibrated ball gap. The aluminum foil is smoothed onto the surface of one of the balls. A capacitor is connected directly across the balls with no current-limiting resistance in series with it. The idea is for all of the stored energy to be delivered to the arc. The amount of energy required to raise a given volume of aluminum to the vaporization point is 10.5 J/mg. Figure 13-3 shows the correlation between calculated and measured hole sizes in the 1-mil foil resulting from different arc energy levels. What results from this test is "calibrated aluminum foil." The efficacy of an actual crowbar system can then be determined by replacing the actual microwave-tube electron gun with a ball gap that has aluminum foil smoothed onto one of the balls and is set to break down at the operating voltage of the actual gun. When this happens, the "let-through" energy can be inferred from the diameter of the resulting hole in the foil. Results of tests of this type have run the gamut from "what hole?" to "what foil?" depending upon the condition of the crowbar system at the time of the test. (Of course, the speed of the subsequent disconnect of the primary power from the high-voltage supply system must be considered also because there can be significant follow-through energy deposition.)

The second component of possible arc damage is the heat generated by ohmic loss in a finite conductor, such as a grid wire. If the arc terminates on it, the resulting "action" of the discharge current can be evaluated as

$$\int i^2 \cdot dt.$$

Any conductor will vaporize, no matter how large in diameter, if it is subjected to sufficient action, or "action-integral," as it is sometimes called. The duration of arc current is short enough that the heating is adiabatic, with no time for heat to be removed. The action required to fuse, or vaporize, copper conductors as a function of their cross-sectional areas is shown in Table 13-1. The units of action are either ampere²-seconds, or joules/ohm. Let's take the circuit of Fig. 13-1 as an

example. If there were no crowbar and the arc terminated on a grid wire, the action could be evaluated as


$$\frac{W_C}{R1 + R2},$$

where W_C is the energy stored in the capacitor C . (This assumes that the total of $R1 + R2$ is much greater than the resistance of the grid-wire segment, which it should be.) If there is a crowbar and it acts so fast that the arc current has not diminished substantially before the crowbar takes it over, the action will be

$$\left(\frac{V_{dc}}{R1 + R2}\right)^2 \times \Delta t,$$

where Δt is the firing delay time of the crowbar. Note that the protective spark gap shown shunting the output of the grid-pulse modulator may succeed in protecting the grid modulator from damage, but it can do nothing to save the grid itself. The only thing that can save the grid is enough surge resistance, enough crowbar response speed, or both. (Of course, the electron gun could be designed so that arcs to the grid itself are all but impossible.)

Grid wires are never made of copper. They are usually made of some metal



Energy (Joules)	Mass (mg)	Area (in. ²)	Hole diameter (in.)
1	0.095	0.00066	0.0288
5	0.476	0.003	0.065
10	0.952	0.0066	0.091
20	1.90	0.013	0.129
25	2.38	0.016	0.144
30	2.86	0.0197	0.158
40	3.81	0.026	0.183
50	4.76	0.033	0.204
60	5.71	0.039	0.224
70	6.67	0.046	0.242
75	7.14	0.049	0.250
100	9.52	0.066	0.289

Figure 13-3. Calculation of aluminum-foil damage due to arc energy.

Table 13-1. Action required for fusing of copper wires of increasing size.

Wire size (AWG)	Area (circular mils)	Area (in. ²)	Fusing action (J/ohm or A ² -s)
44	4.00	3.14x10 ⁻⁶	0.33
42	6.25	4.9x10 ⁻⁶	0.81
39	12.2	9.6x10 ⁻⁶	3.1
36	25	1.96x10 ⁻⁵	12.9
33	50.4	3.96x10 ⁻⁵	52.5
30	100	7.85x10 ⁻⁵	206.6
27	202	1.58x10 ⁻⁴	843.1
24	404	3.17x10 ⁻⁴	3372
21	812	6.38x10 ⁻⁴	13623
18	1620	1.28x10 ⁻³	54233
15	3260	2.56x10 ⁻³	219579
12	6530	5.12x10 ⁻³	881013
9	13090	0.010208	3540261
6	26240	0.02061	14226028
3	52620	0.04133	57208121
1/0	105600	0.08291	230400000
4/0	211600	0.1662	925100000

$$Fusing\ action = \int i^2 dt = \frac{A^2}{33} \log \left(\frac{1083^\circ C - 40^\circ C}{234^\circ C + 40^\circ C} \right) = \left(\frac{Area\ in\ circular\ mils}{6.957} \right)^2$$

that has even lower energy tolerance than copper. When the primary concern is the fusing of internal wirelike conductors, the standard test of crowbar efficacy is to replace the normal high-voltage lead with a length of small-diameter wire (in the AWG 30-to-44 category) and induce an external test arc the same way it was done for the tin-foil crowbar test. Some high-voltage systems have built-in crowbar-test circuits that use a high-voltage relay to short-circuit the output through a length of test wire, greatly facilitating periodic testing. If the wire does not fuse, the crowbar is presumed to be operating properly. When the test wire goes, however, it all goes, and nothing remains between the test-wire terminals. (In a marginal situation, the test wire can actually grow in length over repeated test shots. The wire is being heated to near its melting point while, at the same time, the magnetic force on the wire caused by the fault current is compressing it, making it smaller in diameter. The volume of metal stays the same, however, so the wire must grow in length. Eventually it will vaporize because its diameter is decreasing with each test shot.)

13.2 Types of crowbar switches

The crowbar function can be idealized as the simple, aperiodic, and usually unidirectional discharge of a capacitor bank. The current-limiting surge resistance in series with the discharge path is usually many times the value required for critical damping, $R = 2\sqrt{L/C}$. So any half-control switch with adequate voltage hold-off, peak-current handling capability, and charge-transport lifetime




Figure 13-4. The "infinite-voltage-range" crowbar switch.

can be a candidate. Many successful electronic crowbars have been implemented with hydrogen thyratrons, ignitrons, and triggered air and vacuum spark gaps. However, few, if any, have used solid-state devices, including the most obvious such candidate, the SCR. Because the firing of a crowbar is intermittent with a period between firings that is orders of magnitude greater than the interpulse interval of even the lowest repetition rate periodic systems, it is tempting to make use of the simplest of all discharge switches, the triggered air gap.

A popular form of triggered air gap is schematically shown in Fig. 13-4. It is a simple ball gap with a needlelike mid-plane electrode. The space between the balls is adjustable in order to vary the self-firing voltage. The static voltage on the trigger electrode is maintained at 1/2 of the voltage between the balls by means of a balanced resistive voltage divider. Coupled to the trigger electrode through a pulse-sharpening ball gap is a short trigger pulse that has an open-circuit peak voltage of between 100 and 200 kV, which is obtained from the secondary winding of a step-up pulse transformer. Under normal conditions, the trigger electrode is electrically invisible, lying along the 50% equipotential line between the balls and having the same voltage itself. When the trigger voltage is applied, especially if its polarity is opposite to the system dc voltage, the electrical symmetry of the gap will be instantaneously destroyed, and breakdown between the trigger electrode and the ball connected to the system dc will immediately follow. The channel between the trigger electrode and the other ball will be overvolted by at least a factor of two. The effect will be magnified by the field enhancement at the tip of the trigger electrode. Nevertheless, in a simple, conventional mid-plane-triggered gap, the ratio between maximum self-firing voltage, which must obviously be safely above the maximum system operating voltage, and the minimum system voltage for which there will be reliable breakdown between both balls of the gap is rarely more than 3:1. This means that for operating voltages less than 1/3 of the maximum there will be no reliable crow-

bar protection. This situation is often intolerable.

But note that the gap arrangement shown in Fig. 13-4 is not just a simple mid-plane triggered gap. There are 15- μ H inductors in series with both of the main balls of the gap, and a coupling capacitor, $C1$, in series with the trigger electrode. The purpose of these elements is to transform the simple triggered gap into the "infinite-voltage-range" crowbar switch. The inductors appear as high-impulse impedances to the triggering waveform. When one or the other gap initially breaks down due to the injection of the triggering pulse, the trigger voltage is developed across the inductor in series with the ball terminating the trigger arc. Instead of the trigger electrode being effectively short-circuited to whatever dc voltage is applied to the ball it initially arcs to, the arc pulls the ball up to the trigger voltage, which is momentarily supported by the series inductor, thus allowing the remaining channel to break down as well. Even if the system voltage is such that it initially subtracts from the trigger-pulse amplitude so that the second channel does not break down, the underdamped resonant circuit formed by $C1$ and one or the other of the inductors will cause the trigger-pulse polarity to rapidly reverse, correcting the situation and assuring complete arc-channel breakdown—even with no system voltage applied at all. If it will break down completely with zero system voltage, then its operating ratio is truly infinite, even if the maximum self-breakdown voltage is not.

The version shown is double-ended; it will work with either polarity of system voltage. In most cases, however, the system voltage polarity is fixed and not likely to ever reverse. In this case, a simpler single-ended version can be used that has an inductor in series with the high-voltage connection only and whose trigger-voltage polarity is opposite to the system voltage. To assure operation down to zero system voltage, however, the mid-plane trigger electrode must be mechanically biased slightly toward the ball having the inductor in series with it to make sure that the initial trigger arc is to that ball and not the other one. As system voltage is increased, however, the tendency to break down to the proper ball increases as well.

The presence of inductance in series with the main arc channel does have a price, however. The rate-of-rise of current in the crowbar path will be V/L , where V is the system voltage and L is the 15- μ H inductance. In practical situations, where a little let-through energy is usually a good thing, this is almost never a problem, especially if there is intentional, or even inadvertent, inductance in series with the initial fault-current path. The inductance will limit the rate-of-rise of initial fault current as well.

The performance of a practical application of a single-ended version of this circuit is shown in Fig. 13-5. This crowbar was designed for a dc storage system associated with a high-power dual-klystron RF source. The system's capacitor bank totaled 11.2 μ F. Its maximum voltage rating was 130 kV, although it normally operated at less than 90 kV. The 17.5-ohm surge-limiting resistor was placed in series with the crowbar path, and an additional 21 ohms was placed between the crowbar junction and the klystron cathode bus. Therefore, a total of almost 40 ohms was placed in the klystron gun-arc path. The resistors were of wound construction and therefore had self-inductance, shown as $L1$ and $L3$. The crowbar was tested by simulating an arc path through the crowbar test relay,

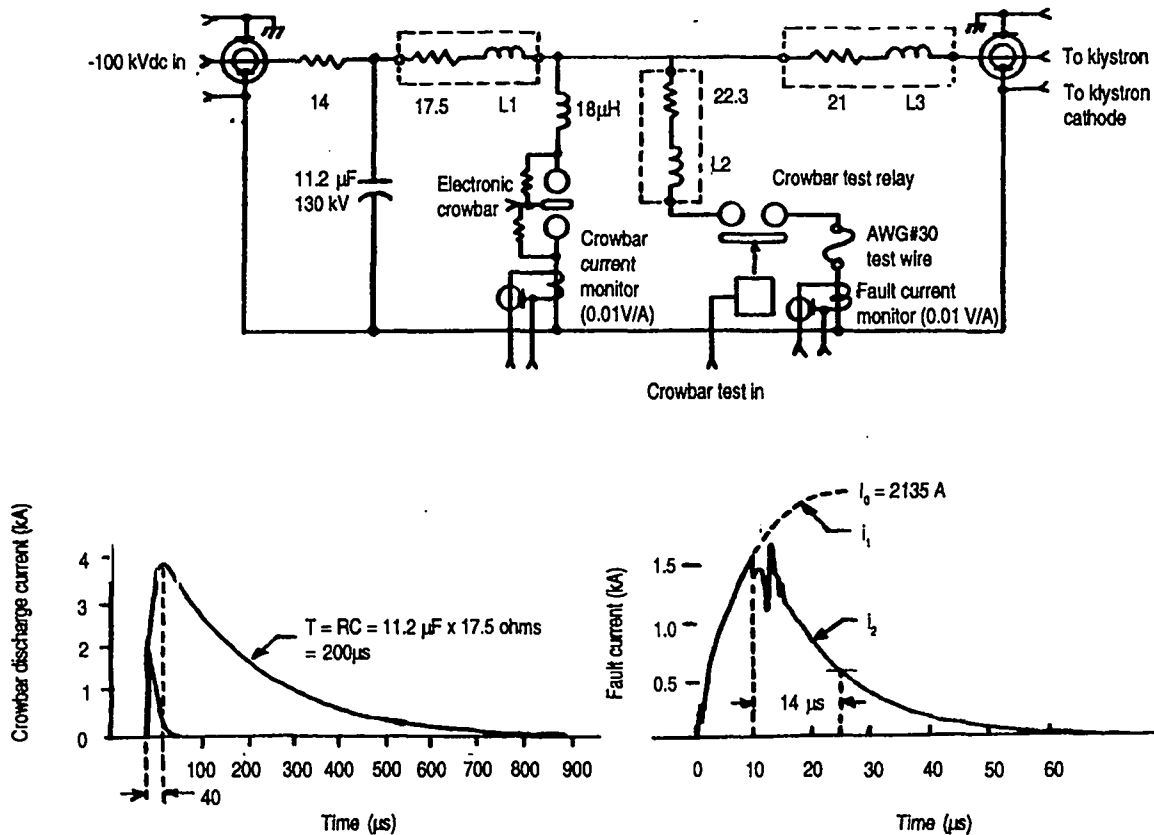


Figure 13-5. Experimental performance of "infinite-voltage-range" crowbar switch.

which had 22 ohms resistance and a length of AWG-30 wire in series with it. Both fault current and crowbar current were monitored by current-viewing transformers.

The waveforms resulting from a test firing are also shown in Fig. 13-5. Fault current in a zero-inductance path would reach a peak of 2135 A at a system voltage of 85 kV, as shown in the following calculation:

$$\frac{V}{R} = \frac{85kV}{17.5\Omega + 22.3\Omega} = 2135A .$$

The actual current waveform is a double-exponential, with a rise time governed by the L/R time-constant and the fall time by the RC time-constant. The crowbar firing delay is 10 μ s, which is neither exceptionally fast nor slow. By the time the gap ionizes, however, fault current, i , has built up to 1500 A, and from that point on it diminishes as discharge current is taken over by the crowbar path. The current through the crowbar continues to build to a peak of 3900 A and then decays with an time-constant of $T = RC = 11.2 \mu F \times 17.5 \text{ ohms}$, or 200 μ s.

Evaluating the fault-current pulse we see that its action, or i^2t time integral, is 26 $A^2\text{-s}$ (or J/ohm), where the integral is made up of the rising current slope, i_1 , and the falling current slope, i_2 . The fusing action for AWG-30 copper wire is 290 $A^2\text{-s}$. No problem there. The total charge transfer is 0.03 C, so the arc energy, assuming 20-V arc drop, is 0.6 J. In all likelihood this is an over-protective

crowbar, one that is not nearly as effective as it could be in the processing of cranky or arc-prone electron guns.

13.3 The nature of surge-limiting resistors

Resistors used in the role of surge-current limiting must be capable of absorbing large amounts of single-shot energy. Those resistors limiting the current through the crowbar must be capable of safely dissipating all of the system stored energy. Those resistors limiting the initial arc current are located between the crowbar junction and the electron gun high-voltage terminal. Although they are removed from the high-dissipation path by the firing of the crowbar, they too must be rated for their fair share of stored energy in case the crowbar does not fire (and often you will not know if it should have fired or not).

These same resistors, especially when used in high-duty-factor applications, must be capable of dissipating large amounts of continuous average power as well, due to the flow of normal load current through them. In a pulsed system this is usually the dissipation rating most likely to be miscalculated. (Remember, it is RMS current not the average current that determines resistor dissipation.) The resistor properties that give it a high-energy rating are not the same as those that give it a high-average-power rating.

What determines the energy absorption capability of a resistor is the mass of the active resistive element and its specific heat, which is about 0.4 J/g-°C for most metals. Remember that we're talking about *active* resistive mass. The mass of the ceramic core of a tape or wire-wound resistor does not count. Only the mass of the tape or the wire itself is important. This is why conventional wire-wound resistors, even with high-average-power ratings, are often poor choices in high-energy applications. There are, however, special high-mass surge windings that can greatly improve their energy handling. Figure 13-6 shows the "ribwound" style of resistor having such a high-mass surge winding. It can often be an excellent compromise where high energy and high continuous-average power must be handled simultaneously. This trade-off is valid because average-power dissipation is determined entirely by the surface area and the maximum operating temperature of the resistive material. The greater the temperature difference between resistive element and ambient environment, the more power a given surface area of resistor can dissipate.

A resistor that is already operating at its average-power-dissipation rating, which is related to a specified temperature rise above ambient (often as high as 375°C), is not in the best condition to take a transient rise in temperature caused by the energy discharge from a capacitor bank. The sum of the temperature rise due to continuous-power dissipation and the transient rise due to specific energy dissipation should ideally not exceed the resistor's maximum specified steady-state rise. The amount of energy that the resistor can absorb for a given temperature rise is given as $W = \Delta T \times \text{mass} \times \text{specific heat}$. One must also remember that average power is nothing more than the time rate of energy. If the capacitor bank is discharged frequently enough, the resistors will never recover to their steady-state temperatures. (Discharging a 1-MJ capacitor bank approximately once every 1.5 minutes is the equivalent of 10-kW average-power dissipation.) As an example of a specific resistor's capability, a ribwound resistor using a 2-kW frame

(3-1/4-in. diameter by 20 in. long) with a 1.45# (700 g) surge winding will absorb 100 kJ with about a 360°C temperature rise. It is obvious that this resistor cannot be used to simultaneously dissipate 2 kW of continuous average power from the flow of normal pulse current at high duty factor and then absorb a 100-kJ shot of energy without being seriously overheated.

The globar style of resistor, also shown in Fig. 13-6, is designed especially for high-energy application. Its entire mass is active resistive material. Its surface-area-to-mass ratio, however, is lower than that of the wound type of resistor, and the thermal conductivity of its material is not as great, so that the temperature gradient in the radial direction cannot be ignored. As an approximate rule of thumb, a globar-style resistor with the same physical dimensions as a ribwound resistor that has a surge winding might have twice the energy-handling capability but only 1/5 the average-power capability. For example, such a resistor that is 1-1/2 in. in diameter and 24 in. long may be rated at 164-kJ energy capability but only 300-W average power dissipation, along with 165-kV voltage hold-off.

Another specially designed type of resistor is the accordion style, also shown in Fig. 13-6. Most of its mass, which is also active resistive material, is shaped like ribbon candy. It can be built in almost any practical length for high voltage hold-off. A resistor rated for 500-kJ energy dissipation, 250-W average power, and 120-kV voltage hold-off has a cross section of 4 in. by 12 in. and is 18 in. long.




Figure 13-6. Different styles of high-energy, high-power resistors.

14. High-Voltage AC/DC Conversion, or Rectification

Vacuum tubes, or electron devices, are unidirectional conductors. Electron charge carriers will only go in one direction. Therefore, at some point in the realization of a microwave-tube transmitter a source of direct current (dc) will have to be provided, and often at very high voltages consisting of tens or even hundreds of kilovolts. The storage of electrical energy in large capacitor banks is also a unidirectional process, even though dc cannot flow through a capacitor. (Charge is stored in a capacitor by the flow of current into it in one direction and removed by a current out of it in the opposite direction. But these directions, once established, do not change.) Almost no microwave-tube transmitters of significant power output are directly powered from a primary source of dc, such as a battery, thermoelectric device, or dc generator. (Some very exotic, single-shot, explosive-energy electrical converters do find uses, though.) Many systems, however, are indirectly powered from such sources, but only after voltage-level conversion that involves the internal generation of alternating current by means of an inverter. Other transmitters are indirectly powered from external sources of alternating current (ac), such as commercial power grids or dedicated alternators. In all of these cases conversion from ac to dc, which is called rectification, is required before it can be used to power our microwave-tube transmitter electron beam (or beams).

14.1 Polyphase ac concepts

Every engineer knows something about ac-dc rectification, but a few still don't know much about polyphase rectification—or even polyphase ac circuits, for that matter. As power levels increase, polyphase circuits become more and more efficacious, and at the higher power levels already discussed they are all but mandatory.

Commercial power is generated and distributed in the form of three balanced phases comprising voltage vectors that are displaced from one another by 120° , as shown in Fig. 14-1. A number of advantages accrue from this arrangement.

- The current in the conductor connecting the neutral, or common point, of a wye-connected source and the common point of a balanced, wye-connected load has a vector sum that is zero.
- The vector sum of voltages around a properly phased, delta-connected source is also zero, allowing the delta to be “closed” without circulating current.
- The vector sum of the mechanical forces on the windings of a three-phase alternator is also zero, so that there are no unbalanced rotational forces.
- The mechanical forces on the rotor of a three-phase motor are purely rotational, which is most desirable.

As we will see, there are additional advantages to three-phase operation when it




Figure 14-1. If you start with two phases, you can make the rest.

comes to rectification.

A balanced, rotating electrical machine is not the only source of a three-phase electrical system, however; nor is there anything magical about just three phases. To create a system having a theoretically unlimited number of discrete phases, all we need to start with is two. And it theoretically does not matter how infinitesimal the phase difference is between them. Even if the difference is a microradian, one vector can be resolved into components that are in-phase and quadrature-phase (I and Q) with respect to the other. Indeed, modern digital signal processing is based on transforming a repetitive waveform into its Fourier series of harmonically related components and evaluating two samples of each, the I and Q components.

Given two voltage vectors of the same amplitude that are in quadrature (which is sometimes called a bi-phase system) and two transformers with identical primary windings and secondary winding ratios and polarities as shown in Fig. 14-

1, a balanced three-phase system can be synthesized. Different transformer winding ratios will accomplish the same thing for different ratios of quadrature-phased vector amplitudes. It also should be obvious that there is no reason to be limited to just three phases. More transformers will produce more phases.

14.2 The three-phase, half-wave rectifier

Single-phase ac rectifiers require external hold-up for either load current (the inductor-input filter) or load voltage (the capacitor-input filter) if the load requires non-varying direct current. Without some form of energy storage between the rectifier and a resistive load, the closest that load voltage and current can ever come to non-varying dc is a succession of half-sine-wave segments, all of the same polarity, all with infinite peak-to-valley ratio. It is direct current, to be sure, but a long way from non-varying. Polyphase ac rectification overcomes this problem.

The simplest form of three-phase rectifier is the half-wave, or three-pulse rectifier, as shown in Fig. 14-2. The transformer secondary is wye-connected. A rectifier diode is placed in series with each phase, and the return from the load is to the neutral point. The rectifier shown is connected for positive-polarity dc output. Whichever phase—A, B, or C—is instantaneously more positive than the other two will be the conducting phase, and conduction will continue for 120° , or $1/3$ cycle, with the peak voltage at the center of the conducting interval. The “natural” commutation point, or transfer point between the phases, occurs when

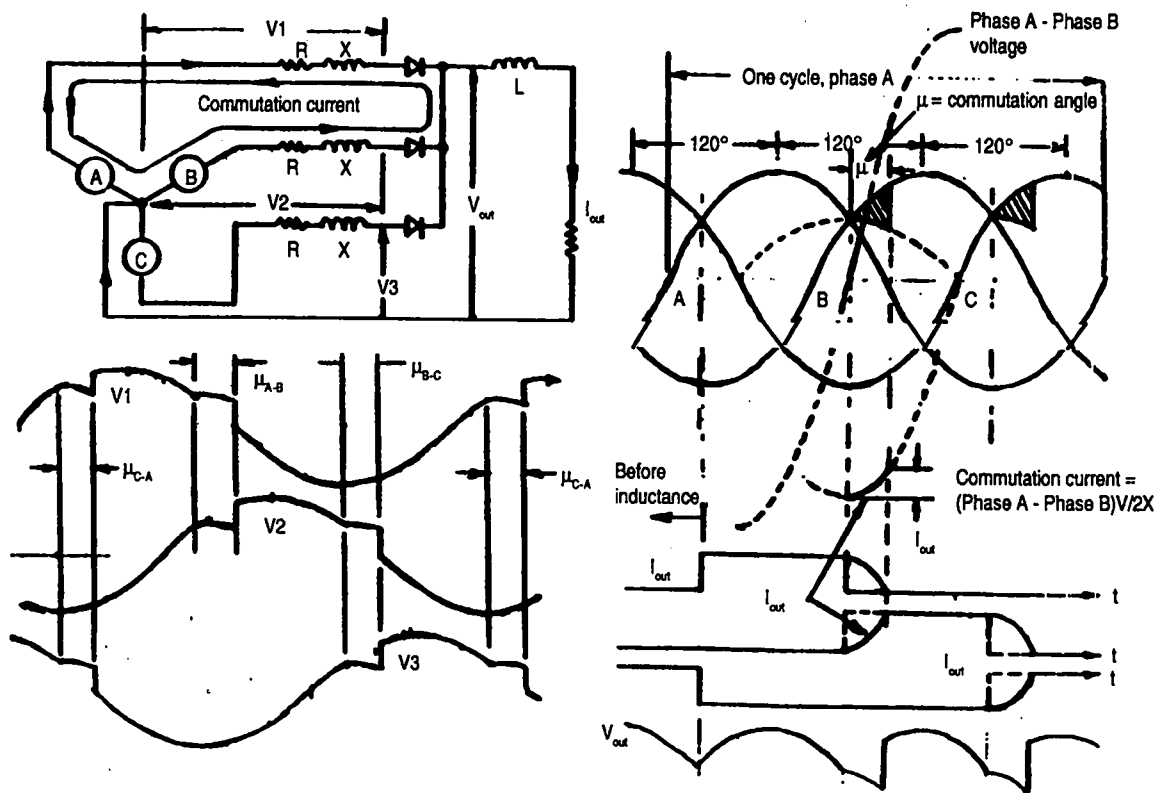


Figure 14-2. The effect of ac line reactance on rectifier current commutation.

the instantaneous voltage of the next sequential phase becomes more positive than the previously conducting phase. At this point for either zero per-phase source inductance or zero load current, the load voltage will simply follow that of the newly conducting phase. This pattern repeats itself throughout successive phase rotations. The load voltage, therefore, will be a series of sinusoid tops spanning each wave's -60° point to its $+60^\circ$ point. There will be three such tops per cycle, hence the term "three-pulse," or "three-bump," rectifier. If the peak voltage is unity, the instantaneous voltage at the commutation point is 0.5, so the peak-to-valley ratio is 2. The theoretical no-load (or no-source-inductance) Fourier series for the frequency spectrum of the load voltage is

$$\frac{3}{\pi} \times \sin \frac{\pi}{3} \left[1 + \frac{2}{8} \cos 3\omega t - \frac{2}{35} \cos 6\omega t + \frac{2}{80} \cos 9\omega t + \frac{2}{(3n)^2 - 1} \cos 3n\omega t \right].$$

The first term has a value of 0.907, which is the ratio of the zero-frequency, or average, "dc" component of output voltage to the peak value of input voltage, even though the valleys of the waveform are only 1/2 as great as the peaks. (Note that the comparable first term for single-phase, full-wave [or two-pulse] rectification is $2/\pi \times \sin(\pi/2)$, which is 0.636. This is the case, even though its order of rectification, or pulse number, is 2/3 that of the simplest three-phase rectifier.) The other terms in the series are the ripple components, the first of which, theoretically, is the third harmonic.

If, however, the three-phase source inductance is not zero—and it never is—and load current is not zero either, the output voltage will not simply follow the tops of the phase voltages. The three-phase source impedance will consist primarily of transformer winding resistance, R , and the reactance of the leakage inductance, X . The resistance is usually no more than 1/5 as great as the reactance and can be ignored as a contributor to total ac impedance. Assuming that inductor L in series with the load is sufficiently large, the output current will be non-varying dc. However, the inductance in series with each phase—inductance through which load current is also flowing—will resist the transfer, or commutation, of the load current to the incoming phase. It will develop whatever voltage across itself that is required to keep its diode forward biased and current going through it. Fortunately, there is a countervailing current called commutation current that opposes the inductor current, eventually forcing it to zero and permitting the commutation of load current to the next phase in line. The commutation current is driven by the vector sum of the outgoing and incoming phase voltages (shown as A and B in the Fig. 14-2) and limited by the reactances of the two phases in series. Heading positive, the driving voltage for commutation current passes through zero at the natural commutation point. Being limited by inductive reactance, however, the commutation current lags the voltage by 90° , hence it is passing through a negative peak at that time. During the time that the commutating current is driving the load current from one phase to the next, which is known as the commutation angle, the output voltage follows neither phase voltage but rather the average of the two. Therefore, the output average voltage is less positive than the incoming phase voltage but more positive than the outgo-

ing phase voltage. At the end of the commutation event, current has transferred to the new phase, and the output voltage abruptly jumps up to coincide with the new phase voltage at that instant. The voltage follows it along until the next commutation point is reached, where the process is again repeated. The relative difficulty of the commutation and the length of the commutation angle are proportional to what is called the commutation factor, which is the product of load current and phase reactance divided by phase voltage.

14.3 The three-phase, full-wave (six-pulse) rectifier

The three-phase, half-wave (or three-pulse) rectifier described above is used primarily in very-low-voltage, high-current applications, where rectifier voltage drop is a primary consideration. Far more popular is the three-phase, full-wave (or six-pulse) rectifier, as shown in Fig. 14-3. Forward conduction of load current occurs for both positive and negative half-cycles of each phase. The no-load output voltage is a succession of 60° conduction intervals, which occur between the natural commutation points from -30° to $+30^\circ$, with respect to each successive peak. There are six such conduction intervals per cycle, hence it is called "six-pulse," or "six-bump," rectification. The valleys are determined by multiplying the cosine of 30° by the voltage peaks, or $0.866 \times V_{pk}$, for a peak-to-valley ratio of 1.15.

The Fourier series describing the theoretical no-load spectrum of the load voltage is

$$\frac{6}{\pi} \times \sin \frac{\pi}{6} \left[1 + \frac{2}{35} \cos 6\omega t - \frac{2}{143} \cos 12\omega t + \frac{2}{323} \cos 18\omega t + \frac{2}{(6n)^2 - 1} \cos 6n\omega t \right],$$

where n is the harmonic number. Ignoring load current and commutation, whose effects on voltage and current waveforms are shown in Fig. 14-3, the no-load spectrum can be broken down into an average dc or zero-frequency term and a series of ripple-frequency components that diminish in amplitude as the harmonic number increases. The value of the first term is 0.995, which is the ratio of the average value of the zero-frequency term to the peak value of the alternating input voltage. The first theoretical ripple component is the sixth harmonic of the input frequency, and its amplitude is $2/35$ of the dc term. Higher-order ripple components are multiples of the sixth harmonic and are diminished by the factor $2/[(6n)^2 - 1]$ from the dc term as the harmonic number increases.

This waveform comes pretty close to what most engineers consider dc, even without additional filtering. In fact, the need for filtering at all is mostly to attenuate the second and fourth harmonic ripple components, which arise from voltage imbalance between the input phase-to-phase voltages. Theoretically—and sometimes in practice—this imbalance can be compensated for directly by separately adjusting phase voltages for greater equality. This procedure will null components that are not multiples of the sixth harmonic in the ripple spectrum. Many high-voltage power-supply vendors will automatically assume there is a three-phase line imbalance of some value—typically 5%—and design filtering to bring the second and fourth harmonic ripple components in line with the overall




Figure 14-3. The three-phase, full-wave, or six-pulse rectifier

ripple requirement.

14.4 The six-phase, full-wave (12-pulse) rectifier

As popular as the six-pulse rectifier has become for high-power, high-voltage applications, the 12-pulse rectifier is rapidly supplanting it as the configuration of choice. Even though it requires two full-wave bridge rectifiers instead of one and two separate secondary windings to feed them, the total number of rectifier junctions is usually the same. This is because series-connected individual rectifier junctions are required to provide the total reverse-voltage capability and it does not matter whether they are connected as one bridge or two. The transformer core and coil combination is also much the same, except the secondary is divided two halves from a total volt-ampere perspective.

The six-phase input required for a 12-pulse rectifier can be derived in several ways. A "natural" way is by the use of a dual secondary, one winding of which is connected in delta, and the other in a wye. Not surprisingly this design is referred to as a "delta-wye" configuration and is shown in Fig. 14-4. A disadvantage of this arrangement is that the two secondaries are not identical, except in volt-ampere product. The legs of the delta winding must have $\sqrt{3}$ times the voltage of the wye-connected legs. To even things out, the wye-connected legs must handle $\sqrt{3}$ times the current. The wye windings have fewer turns but are of heavier gauge wire than the delta-connected windings. This makes things a bit more complicated for the winding shop (and therefore makes it easier to get something wrong). The "natural" aspect of the six-phase output is actually illustrated in the figure. Note that if the windings were drawn to a scale proportional to their voltages and aligned so that they were parallel to their corresponding primary winding, the vectors representing the resultant, or line-line, voltages of the wye-connected winding would be tilted by 30° with respect to the corresponding delta-connected output vectors. This 30° "interphase" displacement interleaves the outputs of the two six-pulse rectifiers in series, yielding a 12-pulse output with conduction angles of 30° for each output pulse, or "bump." The phase progression is shown in Fig. 14-5, where each phase—A, B, and C—has a "shadow" displaced by 30° —A', B', and C'.

Except for the effects of any three-phase line-voltage amplitude imbalance,




Figure 14-4. The "delta-wye," six-phase, full wave, or 12-pulse rectifier.




Figure 14-5. The phase progression for 12-pulse rectification.

the six-phase, 12-pulse rectifier produces nearly pure non-varying dc. Disregarding the commutation effect of source inductance, the valleys in the time-domain output voltage are 97% of the peaks, for a peak-to-valley ratio of 1.035. The Fourier series describing the theoretical no-load output spectrum is

$$\frac{12}{\pi} \times \sin \frac{\pi}{12} \left[1 + \frac{2}{143} \cos 12\omega t - \frac{2}{575} \cos 24\omega t + \frac{2}{1295} \cos 36\omega t K + \frac{2}{(12n)^2 - 1} \cos 12n\omega t \right].$$

The first term is 0.9886. The first unavoidable ripple component is the 12th harmonic, which is already attenuated by the factor $2/143$, making it 37 dB below the "dc" term. It should be noted that high-voltage power supplies for use in applications that require extremely low ripple amplitude (such as -120 dB) at relatively high ripple frequencies (such as 10 kHz or above) often use rectifiers like this operated from 60 Hz. A component at 10.08 kHz, for instance, would be the 168th harmonic of 60 Hz, having a harmonic number, n , of 14. It would be attenuated by a factor of $2/[(12 \times 14)^2 - 1]$, or $1/28223$, which is -89 dB without filtering. The residue left after filtering is often the noise produced by corona, or partial-discharge, resulting from imperfect high-voltage design. Such a corona often becomes the irreducible noise floor.

The disadvantage of the "delta-wye" transformer secondary arrangement is

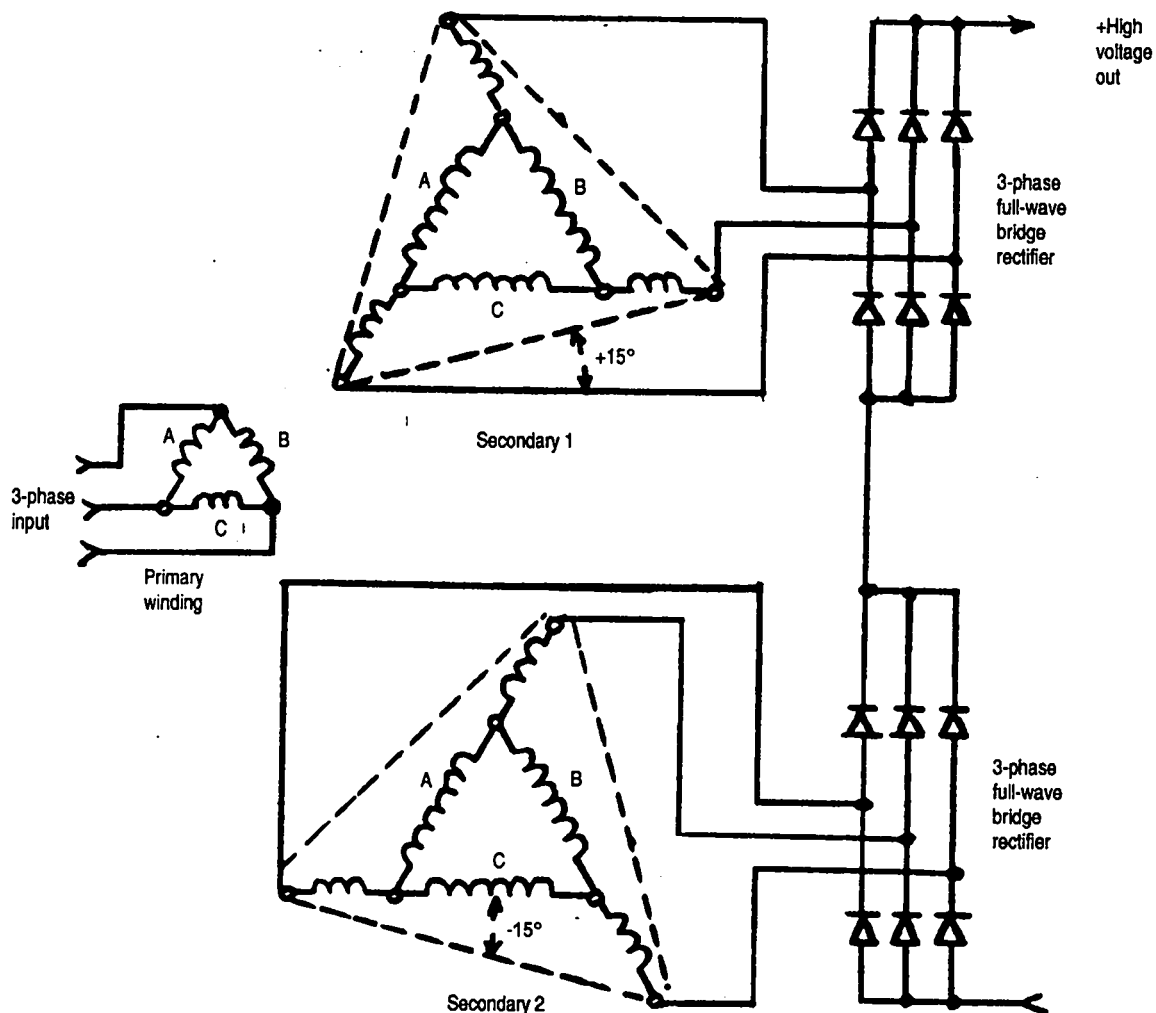


Figure 14-6. The "extended-delta" six-phase, full-wave, or 12-pulse rectifier.

largely overcome by the "extended-delta" pair of secondary windings, as shown in Fig. 14-6. In this case, identical secondary windings can be used. The windings for each phase have taps, or "extensions," at both ends. If the taps are properly placed, making the connections shown will produce a $+15^\circ$ "tilt" to the line-line resultant voltage vectors of the upper secondary and a -15° tilt to those of the lower one, thus generating the 30° interphase angle between the outputs of the series-connected full-wave rectifiers.

14.5 The 12-phase, full-wave, 24-pulse rectifier

Figure 14-7 shows an extension of the 12-pulse rectifier into a 24-pulse rectifier. This is an actual commercial design, comprising a pair of complete "delta-wye"-type 12-pulse rectifiers, each with its own complete transformer. The primary windings of the two transformers are identically wound and connected in a configuration called "almost delta," or "polygon." For both primary windings, a small C-phase winding segment is connected between the A and B junction, a small A-phase winding segment is connected between B and C, and a small B-phase winding segment is connected between C and A. Therefore, all phases are identical. However, the feed lines from the incoming three-phase source are not

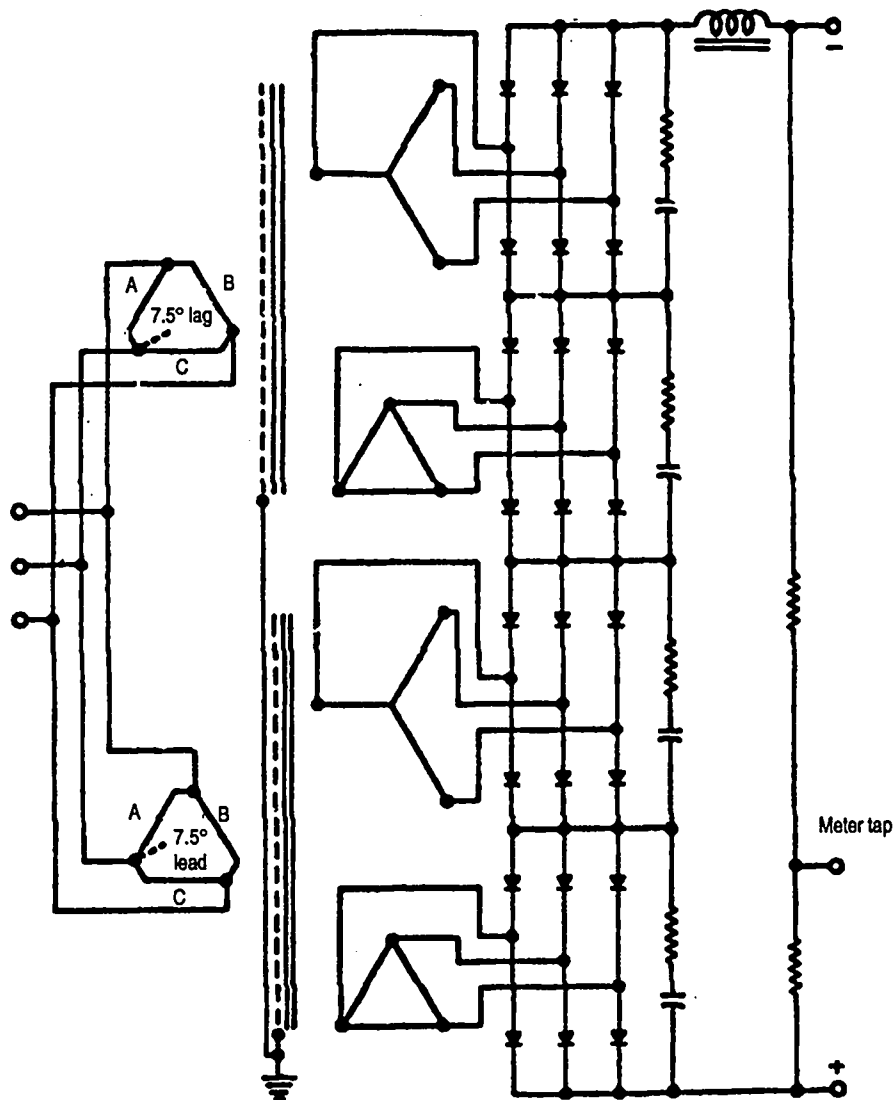


Figure 14-7. A 24-pulse rectifier system that uses polygon primaries.

connected to identical points. The difference imparts a 7.5° lead to one resultant voltage vector and 7.5° lag to the other for a total difference of 15° "inter-inter-phase." In other words, that the two 12-pulse rectifier outputs are interleaved to give 15° conduction angle for each sequential phase segment.

As can well be imagined, there is no theoretical upper limit to how far such polyphase upscaling can be carried, but there is certainly a point of diminishing returns. Although more than a few 24-pulse-rectifier systems are currently operating, this configuration is considered the practical upper limit.

14.6 Polyphase-rectifier line current

So far, the discussion of polyphase rectification has dwelt upon design and user-friendly aspects, such as how close can the outputs come to true non-varying dc, and how tenuous is the ripple spectrum. There is another reason to admire polyphase rectification that is more selfless and environmentally friendly: it inherently ameliorates line-current harmonic pollution.

Consider the typical front end of an “off-line” switch-mode ac-dc converter operating from single-phase ac input, as shown in Fig. 14-8. It is designed to generate a 280-Vdc “rail” from either 115-Vac or 230-Vac input by using the rectifiers as either a full-wave bridge (230 V) or as a full-wave voltage doubler (115 V), depending upon the switch position. The output voltage from the rectifiers into a resistive load bears little resemblance to non-varying dc. For this reason, huge hold-up storage capacitors, C_H , are charged to the peak value of the incoming alternating voltage and hold the value nearly constant until the next positive peak of full-wave rectifier output. The rectifier conduction angle, however, tends toward zero because the rectifiers are back-biased for most of the time.

All of the power consumed is contained in the product of the line voltage and the fundamental-frequency component of line current, which is in phase with the voltage. The total RMS value of the spikelike input current waveform can be two to three times as great as the in-phase, fundamental-frequency component, and the non-useful contributors to the total are virtually all harmonic components because the fundamental-frequency component is very nearly in phase with the line voltage. (When the fundamental-frequency component of line current is not in phase with the voltage, the effect on power factor is called “displacement,”

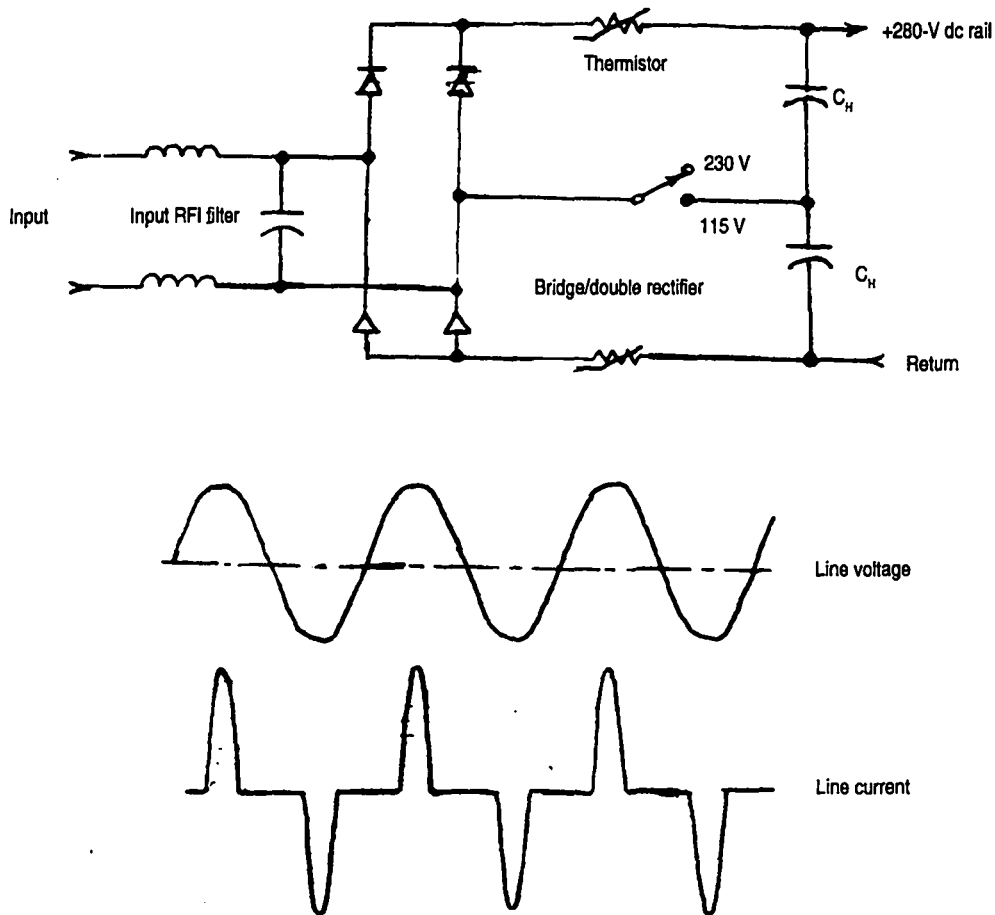


Figure 14-8. Typical front end of off-line, single-phase power supply and typical current wave shape.

and it is proportional to the cosine of the relative phase. However, when the decreased power factor results from harmonic components of line current, the effect is described as "distortion.")

Even though the conversion efficiency of such a supply can be over 95%—the only losses being diode-voltage drops—the power factor, which is the ratio of the input power to the total input volt-amperes (including harmonic components), can easily be less than 0.5. This means that the total RMS input current can be twice as great as the useful component, thus heating up feed lines and circuit breakers four times as much as the useful component would by itself.

The power factor of such a supply would be improved if the dc filter input were inductive instead of capacitive. If the inductor were large enough, the load current would be non-varying dc, and the line current would then be a square wave that has a peak amplitude equal to the dc output current. The peak value of fundamental-frequency component that would be $4/\pi$ times the square-wave peak amplitude and it would be in phase with the voltage. The RMS value of this current is

$$\frac{4}{\pi} \times \frac{1}{\sqrt{2}} \times I_{pk},$$

and the RMS value of total input current is I_{pk} . Assuming the line voltage has a peak value of V_p , the average dc output voltage would be $2/\pi \times V_p$, and the RMS value of input voltage would be

$$\frac{V_p}{\sqrt{2}}.$$

The dc output power would be

$$V_{dc} \times I_{dc} = \frac{2}{\pi} V_p \times I_{pk}.$$

With no rectifier or inductor losses, the input and output power would be equal. The input power is the product of the in-phase fundamental-frequency components of line voltage and current, which is

$$\frac{V_p}{\sqrt{2}} \times I_{pk} \times \frac{4}{\pi\sqrt{2}} = \frac{2}{\pi} V_p \times I_{pk},$$

as shown above. The total input volt-ampere rating is

$$\frac{V_p}{\sqrt{2}} \times I_{pk}.$$

The ratio of the two is the power factor, which in this case is 0.9—a big improvement over the capacitive input filter. However, there is a problem with inductor

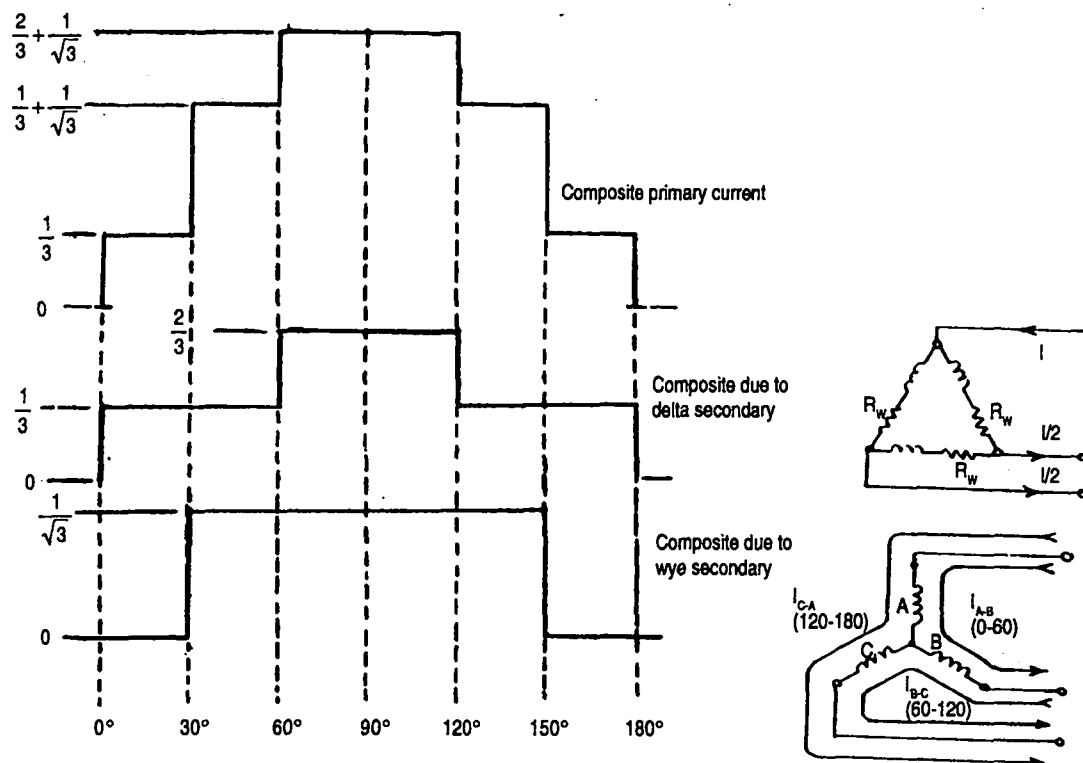


Figure 14-9. Theoretical input line current for delta-wye 12-pulse rectifier.

inputs.

Inductors of the type that will do the job are large, heavy, and expensive, and they subvert the purpose of the high-frequency switch-mode converter that follows the off-line rectifier. In addition, the rail voltage will be lower in the bridge-rectifier mode, and the voltage-doubler will not work at all.

Consider, instead, the waveforms for the line current resulting from the delta and wye windings of a 12-pulse rectifier, shown in Fig. 14-9. With a resistive load on the rectifier, the load current will be substantially constant even without a large input inductor filter, because the rectifier output is so close to non-varying dc to begin with.

The component due to the delta winding will have a peak value of $2/3$ of the output current, which flows for the 60° interval centered on the peak. For the other 120° of each half-cycle, the current will be $1/3$ of the peak. This is because the current, I , will split at the delta junctions in accordance with the dc resistance of the windings. Each phase winding is shunted by the other two in series, thus this configuration has twice the dc resistance of the single-phase winding. Therefore, the current will split with the ratio of 2:1, so $2/3$ will flow in the phase winding having the highest instantaneous voltage and $1/3$ will flow in the series combination of the other two.

The component due to the wye-connected winding will have a conduction angle of the center 120° of each half-cycle because, at any given time, current is flowing in 2 of the 3 phases. Each phase is idle only $1/3$ of the time. The amplitude of the current will be $1/\sqrt{3}$ as great as the dc value because of the $\sqrt{3}$ ratio of the delta and wye turn ratios to ensure equal rectifier input voltages.

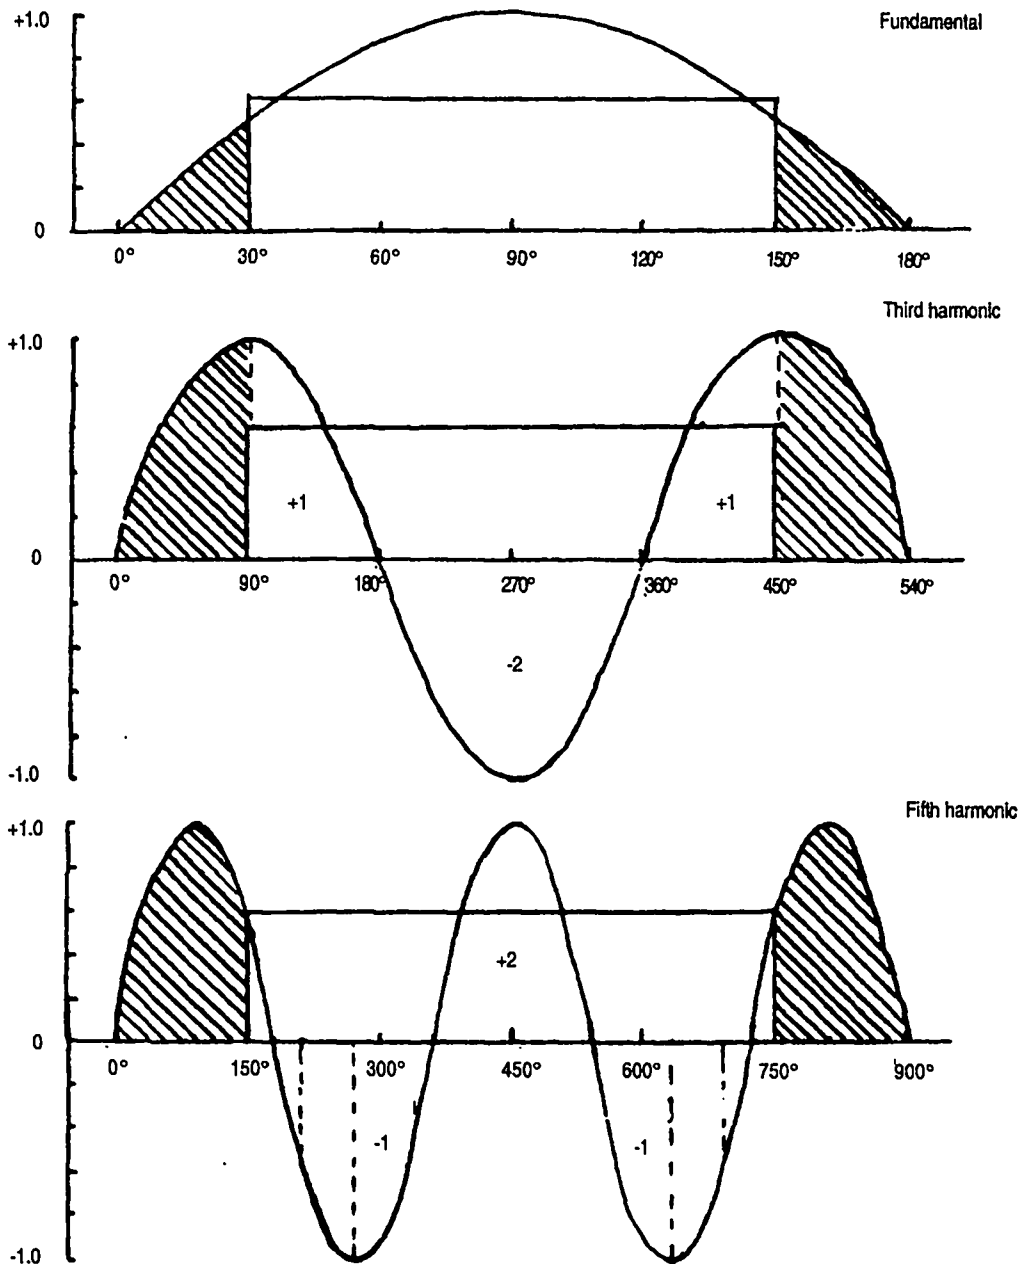


Figure 14-10. Graphical-analysis Fourier series through the fifth harmonic of a six-pulse rectifier line current.

The theoretical Fourier series for the line current of a rectifier system having a pulse number of P is

$$\begin{aligned}
 I = & \sin \omega t + \frac{1}{P-1} \sin(P-1)\omega t + \frac{1}{P+1} \sin(P+1)\omega t \\
 & + \frac{1}{2P-1} \sin(2P-1)\omega t + \frac{1}{2P+1} \sin(2P+1)\omega t \\
 & + \frac{1}{nP-1} \sin(nP-1)\omega t + \frac{1}{nP+1} \sin(nP+1)\omega t .
 \end{aligned}$$

Note that the line-current harmonic components are the odd values just below and above the rectifier-output harmonic components. For instance, a six-pulse rectifier would have a theoretical first ripple component of the sixth harmonic, while the first line-current harmonics would be the fifth and seventh. For those who might wish to see a very simple-minded proof that requires no mathematics at all, Fig. 14-10 illustrates this phenomenon for the line current due to the wye-connected secondary winding. If we disregard the coefficients and concentrate on the harmonic-component ratios, a Fourier series is evaluated by multiplying the subject waveform by all of the harmonically related sine and cosine waveforms (in-phase and quadrature). Assigning a value of unity to the line current, the fundamental can be seen to be a sine-wave segment from 30° to 150° . The third harmonic sine wave is multiplied by zero between 0° and 90° and again between 450° and 540° . For the remainder, just using the similar shapes, there are two positive quarter-cycles (+1) and one negative half-cycle (-2), which cancel, so there is no 3rd harmonic component. (If we trace out the cosine waveforms, we can easily see that they all cancel to zero.) The fifth harmonic sine wave, however, has non-zero product between 150° and 750° . In the center, two negative quarter-cycles cancel the one positive half-cycle and the two positive and negative 30° segments at the two ends cancel each other. What is left is two negative segments, one from 210° to 270° and the other from 630° to 690° . If they are pushed together, they have the same wave shape as the fundamental component, except they take up $1/5$ of the area.

If one does the same "mathless" analysis of the primary current due to the delta-connected secondary, the same wave shape emerges for the fifth harmonic. But it is positive instead of negative, so it would cancel the fifth harmonic component due to the wye-connected secondary in the 12-pulse rectifier. As the pulse number of the rectifier is increased, it can be seen that the line current becomes less and less polluted by harmonics and increasingly resembles in shape a fundamental-frequency sine wave.

14.7 Who cares about harmonic pollution?

The electrical power utilities certainly care about harmonic pollution. In the first place, they do not get paid for the current it represents, but they do have to suffer the power loss it causes. Furthermore, "trapped" resonances at the harmonic frequencies can cause all manner of electrical mischief in terms of resonant overvoltage. And from almost from the beginning of the electrical age, the ability of power-line harmonics to affect other electrical services has been noted.

The telephone company, whose wires parallel those of the power company in




Figure 14-11. Graphs showing changing sensitivity of telephone equipment to coupled harmonic signals.

many situations, has been aware for a long time of the "metallic" noise that power-frequency harmonics can induce. Figure 14-11 shows some graphs that trace the evolution of the most troublesome frequencies the telephone company has had to endure as its equipment has improved in frequency response and fidelity. The graph plots the disturbance frequency against the Telephone Influence Factor (TIF), which is a number that defines relative interference. In 1919, for instance, the most offensive frequency component was barely 1 kHz. This did not mean that the phone equipment itself was the most sensitive at that frequency, because the TIF includes a mechanism that takes into account the coupling between power and phone wires. This coupling increases almost linearly with frequency. By 1960, a broader range of frequencies, peaking at 2800 Hz, became troublesome. Today, harmonic pollution can also disrupt digital data sent from modem to modem. (But the TIF concerns will soon fall by the wayside when all telephone signals are transmitted by fiber-optic cables.) To reduce this pollution problem, some governments and organizations have promulgated strict standards. In fact, present European standards for line-current harmonic content, as well as those standards prescribed by the U.S. Navy, are far more stringent than even the phone company could wish for.

Table 14-1, which tabulates the odd harmonics that would appear in the line current of a six-pulse rectifier ($n \times 6 \pm 1$) through the 73rd harmonic, shows the theoretical and actually measured values of harmonic current for a six-pulse rectifier. It then compares them with a 24-pulse rectifier that the six-pulse rectifier is part of. The values of current are then multiplied by the TIF weighting factor (T) for that particular frequency, which is shown in the third column from

the left. The composite TIF, which is only a relative term, is obtained by taking the root of the sum of the squares of each $I \times T$ product and dividing that by the fundamental-frequency current. Note that the six-pulse rectifier has a TIF of 277, whereas the 24-pulse-rectifier TIF is only 60, which is a marked improvement.

14.8 Other forms of line pollution

Another form of objectionable line-current disruption is "flicker" modulation, which is caused by load changes that have a relatively low recurrence rate. Although this type of modulation is by no means peculiar to rectifier loads, the electronic agility of most systems fed by high-power rectifiers make them the most likely to produce this sort of pollution.

The "flicker" referred to is the modulation in the light output of an electrical lighting system. It is most often seen in incandescent light bulbs. Often, large electronic systems will actually have a flicker-modulation specification applied to them that defines a maximum change in line voltage they can produce. Equally often, the specification will be misapplied. How much the line voltage changes depends not only on the variational current demand of the system but on the internal impedance of the power source. An extremely stiff, or low-impedance, source will not permit voltage variation, no matter what the line current does.

Table 14-1. Weighted Telephone Influence Factor calculations for six-pulse and 24-pulse rectifier systems.

Harmonic number	Frequency	TIF Value	Six-Pulse Rectifier			24-Pulse Rectifier		
			Measured line current	Theoretical line current	IxT product	Measured line current	Theoretical line current	IxT product
Fund	60 Hz	0.5	74 A	74 A	37	604 A	604 A	302
5	300	225	16.7	14.8	3760	2.42	0	540
7	420	650	6.47	10.6	4210	1.1	0	720
11	660	2260	4.6	6.7	10,400	0.5	0	1200
13	780	3360	2.98	5.7	10,000	0.33	0	1100
17	1020	5100	1.56	4.3	7960	0.06	0	260
19	1140	5630	1.17	3.9	6590	0.22	0	1220
23	1380	6370	0.5	3.2	3100	4.1	26.3	26,100
25	1500	6680	0.34	3.0	2800	3.2	24.2	21,300
29	1740	7320	0.46	2.6	3400	0.17	0	1240
31	1860	7820	0.31	2.4	2400	0.07	0	560
36	2100	8830	0.44	2.1	3900	0.14	0	1210
37	2220	9330	0.33	2.0	3030	0.15	0	1380
41	2460	10340	0.24	1.8	2480	0.04	0	450
43	2580	10600	0.19	1.7	2050	0.09	0	960
47	2820	10210	0.12	1.6	1230	0.95	13	9660
49	2940	9820	0.08	1.5	790	0.76	12.3	7500
53	3180	8740	0.14	1.4	1190	0.06	0	550
56	3300	8090	0.8	1.35	660	0.08	0	650
59	3540	6370	0.14	1.25	930	0.14	0	930
61	3660	6130	0.1	1.2	630	0.06	0	360
65	3900	4400	--	1.1	--	0.03	0	130
67	4020	3700	--	1.1	--	0.05	0	190
71	4260	2750	0.07	1.0	180	0.41	8.5	1100
73	4380	2190	0.03	1.0	70	0.37	8.3	810

Root sum square of 6-pulse $I \times T$ product = 20,500
Weighted Telephone Influence Factor = 20,500/74 A = 277

Root sum square of 24-pulse $I \times T$ product = 36,100
Weighted Telephone Influence Factor = 36,100/604 A = 60

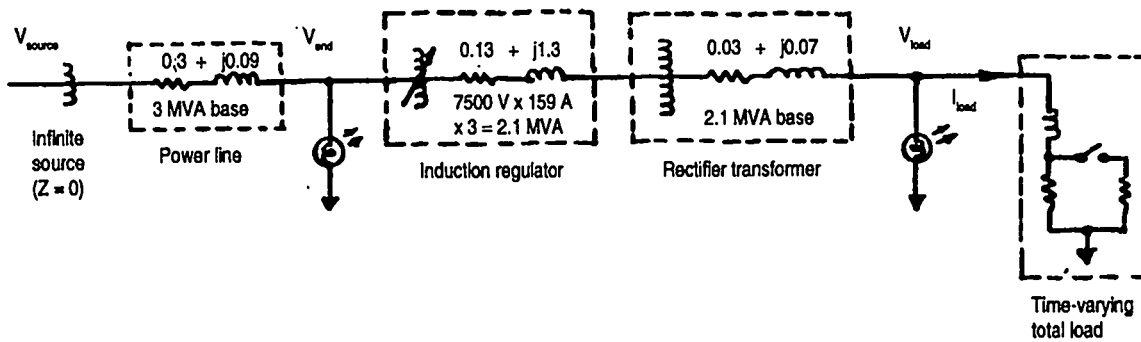


Figure 14-12. Typical line impedance for large transmitter system.

Not all sources are stiff enough, however, to keep flicker within prescribed limits. Therefore, they may need to be investigated.

Figure 14-12 shows the impedance diagram of an actual high-power feed, in this case the one that supplies the Haystack Hill complex. The reason that it was of interest is that the LRIR had as one of its modes of transmitter operation a 50% duty factor at a 10 pps repetition rate. The peak demand during the 50-ms "on" time, was more than 2 MW. It was assumed that picking up and dropping a 2-MW load 10 times per second might not go unnoticed. Reviewing Fig. 14-12, therefore, will provide a little refresher in normalized power-source impedance and the concept of VA base, or, in this case, MVA base.

The eventual MVA base, 2.1 MVA, is that of the transmitter dc power supply. To determine the total source impedance in per-unit or percent, it is necessary to normalize the other components to the same base. The resistive components of source impedance can be ignored because their contribution to total impedance is negligible. The significant impedance of the rectifier transformer, $j 0.07$, is almost entirely leakage reactance and is already normalized to its own load-impedance base of 2.1 MVA. The reactance of the variable transformer ahead of it, the induction regulator, is given in absolute terms, as $j 1.3$ ohms. To normalize it we must find its load impedance, which is the square of its maximum output voltage, 7500 V, divided by its MVA base, 2.1 MVA, or 26.8 ohms. Therefore, the per-unit value is therefore $j 1.3$ ohms divided by 26.8 ohms, or $j 0.05$. What remains to consider is the power line that feeds the system from the virtually infinite source. Its per-unit reactance is $j 0.09$, but its VA base is 3 MVA instead of 2.1 MVA. To convert this value it is necessary to multiply the reactance by $2.1/3$, giving $j 0.063$ as the per-unit reactance on a 2.1 MVA base. The total reactance




Figure 14-13. Voltage regulation of power system shown in Fig. 14-12.




Figure 14-14. Subjective properties of "flicker" modulation.

then is $j(0.063 + 0.05 + 0.07) = j 0.18$, or 18% reactance, which is a long way from being "stiff."

Figure 14-13 shows how the load voltage at full-load would differ from the source or no-load voltage. The load power factor is approximately 0.8, so the line current will lag the load voltage by 37° . The voltage drop across the source reactance will be in quadrature with the current. The vector sum of source voltage and reactance voltage drop is 0.895 times the source voltage, indicating about 10% no-load-to-full-load voltage change. Is this a little or a lot?

Figure 14-14 shows two curves that are the results of subjective responses of people to the lighting flicker caused by different voltage changes at different flicker rates. One curve plots the borderline of irritability, the other the borderline of visibility. The total area above each curve represents an area where the values of voltage change produce visible flicker or induce irritability. Note that humans beings are most sensitive, both in terms of visual perception and emotional irritability, when the flicker rate falls between 8 and 10 fluctuations per second (sound familiar?). A voltage fluctuation of about 0.25% produces observable lighting flicker, whereas voltage fluctuation of less than 0.5% brings on irritation. Needless to say, the high-voltage dc-power-conditioning system for the Haystack Hill LRIR transmitter required serious load-transient mitigation in the form of a large energy-storage capacitor bank, electronic voltage regulation, and current regulation to keep the load on the power line constant regardless of operating mode on the dc output side.

15. Variable-Voltage Devices

If you ever design a high-voltage dc system for a microwave tube, make sure it has a mechanism for varying the voltage, preferably from zero on up. This is true even for microwave tubes like TWTs and crossed-field devices, which have narrow ranges of actual operating voltages. The commissioning of almost every component in the high-voltage loop is a gradual affair, and it is all but impossible to commission a system without being able to control the high voltage over nearly its complete range. It is also important that the variability be low-loss because we are usually dealing with large amounts of average power.

15.1 The phase-controlled rectifier

One way to achieve variable high voltage is by the use of phase-controlled rectifiers in the ac-dc converter. Indeed, this is the basic application of the SCR, which has already been discussed. Before the advent of solid-state components, however, devices such as thyratrons and ignitrons were also used as phase-controlled rectifiers in variable-output high-voltage supplies. A full-control rectifier system, illustrated in Fig. 15-1 as a six-pulse rectifier, requires nothing more complicated than the replacement of the rectifier diodes with SCRs. A half-control rectifier involves the replacement of only half of the rectifiers, but a complication arises in the gate-firing circuits.

The electronic strategy used to produce less than the full-output voltage for a given ac input is to delay the gating of the controllable rectifiers until after the normal commutation point (in time) has passed, as shown in Fig. 15-2. The conduction angles of the six "CRs"—they do not have to be silicon—are the same, but they are less than the normal 60° intervals that produce maximum-output voltage. The waveforms to the left in the figure would apply under no-load conditions for different retardation angles, which are the amounts by which the rectifier turn-on points are delayed from the "natural" commutation points. On the right side of the figure are the waveforms that are obtained under load




Figure 15-1. The phase-controlled rectifier variable-output dc power supply.

and with finite three-phase source inductance, giving rise to a commutation angle as well as a retardation angle. Note that the output-voltage waveform with retardation includes an instantaneous jump upward halfway between the outgoing- and incoming-phase voltages. This is followed by downward-tracking segment that is also midway between the phase voltages. Finally, when the stored energy in the source inductance has been overcome, another step upwards occurs to meet the instantaneous voltage of the incoming phase. This complex voltage waveform has a lower average value than the full-conduction waveform. It also requires a rectifier load having a high-variational input impedance, such as an inductor-input filter, in order to support the voltage.

The ratio of partial-conduction output to full-conduction output is the cosine of the angle of retardation (ignoring the effect of commutation delay). For the case of 60° ($\pi/3$ radians) retardation, the waveform looks like a sine wave traveling between 0 and $\pi/3$ radians, but backwards. The integral of this waveform is $1/2$. Its average value is $3/\pi \times 1/2$. The full-voltage output has an average value, as previously discussed, that is $6/\pi \times \sin(\pi/6)$. But the sine of $\pi/6$ is $1/2$. This gives $6/\pi \times 1/2$, which is twice as great as the average value for 60° retardation. We would expect this result because the cosine of 60° is $1/2$. For 90° retardation, the waveform at the input to the inductor is finite, comprising positive and negative 30° segments. But their average value is zero, as would be expected




Figure 15-2. Phase-controlled rectifier waveforms for different output voltages.

because the cosine of 90° is 0.

Note that this is a true switch-mode application of an inherently lossless device, the half-control switch. In order for there to be dc output, the rectifiers must be gated "on" during each and every conduction interval. They will block forward and reverse voltage unless they are gated into conduction in the forward direction. Although this arrangement implies that there be considerably more electronic complexity than for a more conventional dc power supply, it also provides the mechanism for the highest-speed load disconnect following fault conditions: simply blocking the low-level SCR gating signals. Gate-signal timing strategies can also produce such desirable features as voltage or current regulation, soft-start (or ramped-up turn-on), and output current-limiting.

High-voltage variable-output power supplies using phase-controlled rectifiers are relatively rare, but they are by no means unheard of. A new installation, part of a test facility for the microwave-tube group of the French corporation Thomson, uses phase-controlled SCRs to vary the output of a 30-kV dc power supply to charge pulse-forming networks in a highly versatile line-type modulator. (This installation was designed by the same engineers who produced France's high-speed train. Experience helps.)

15.2 The primary SCR controller

A far more popular application of the half-control switch in the role of a variable-voltage device is as the active element of the primary-voltage controller. Because the SCR is the device that predominates in this area—although it is by no means the only one—the primary-voltage controller is universally referred to as the primary SCR controller. It differs from the phase-controlled rectifier in that what is varied is the amplitude of the alternating voltage applied to the primary winding of the rectifier transformer. It works in much the same fashion as the

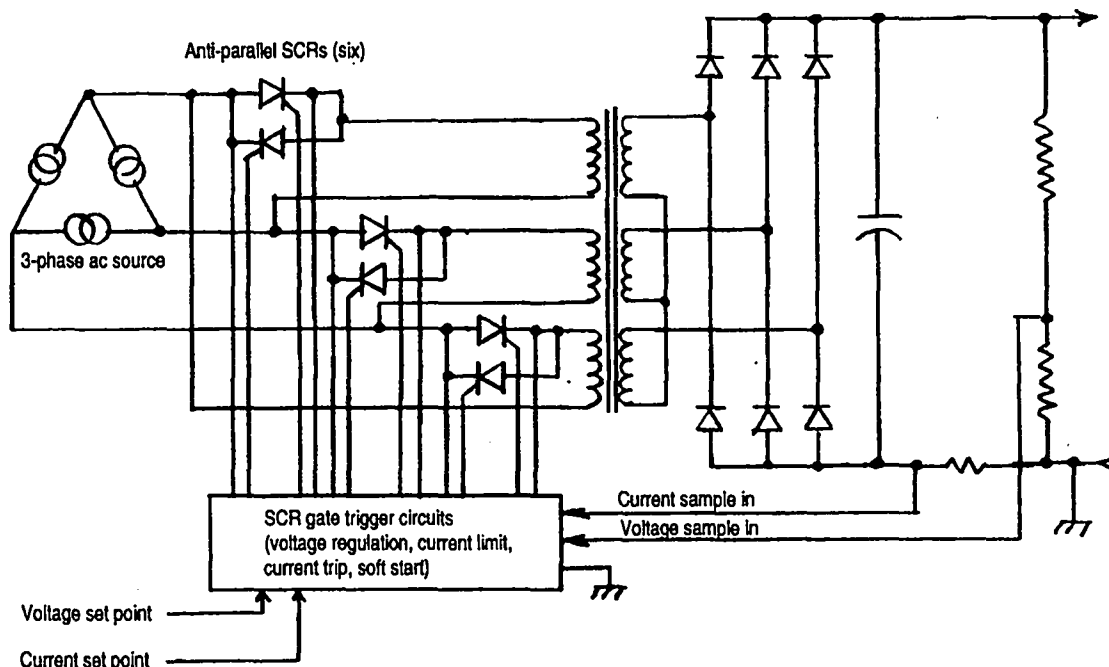


Figure 15-3. Primary phase control of variable dc output voltage using SCRs.

household full-wave light-bulb dimmer. Of course, primary current is bidirectional, which is why the control devices, as shown in Fig. 15-3, must be provided in pairs connected in anti-parallel configuration (parallel, but of opposite polarities).

At first glance it is apparent that many more devices are now required to produce the same voltage variability of the phase-controlled rectifier. Not only are a minimum of six primary half-control devices required, but rectifiers on the secondary side must be provided as well. The great advantage of this design is that the primary voltage can be selected to match the voltage-handling capability of a single device, even if it must often have prodigious current-handling ratings. Primary voltages rarely exceed 600 Vac, even for power supplies rated in megawatts of dc output and hundreds of kilovolts of output voltage. The gate-control circuits, although not simple, are by no means as complicated as the comparable gating circuits required by series-connected SCRs in a high-voltage phase-controlled-rectifier supply. Integrated low-level controls, comprising voltage and/or current regulators and "proportional" SCR gating circuits, can accomplish all of the same functions described earlier. (The gating circuits are called proportional because the effective primary voltage is proportional to the low-level dc input signal to the trigger board.)

High-speed disconnect from the ac source and current-limiting into a short-circuited load are important features of proper gate-control strategy. These are not the only design considerations, however. SCRs can and do fail, especially when they are subjected to transient source overvoltage, such as from a lightning strike (or near miss), power-line switching transients, etc. When SCRs fail, they take on all of the characteristics of bus bars. So transient suppression in the source is not an option; it is a necessity. Also, the SCR controller must always be backed up by an air-insulated contactor. Even though it is slow-acting, it can ensure positive disconnect from the ac source. SCR control on the primary side is much more susceptible to loss of control than phase-controlled rectification. Ironically, this is partly due to one of the phase-controlled-rectifier's disadvantages when used at high-voltage: many rectifiers may be needed in series to handle the full voltage. In such circuits, the failure of one or two phase-controlled rectifiers may not cause immediate failure of the entire series string. (In fact, periodic checks of series strings of conventional high-voltage solid-state rectifiers regularly indicate that as many as 10% of the individual junctions are short-circuited at any given time.) But the failure of one or more SCR primary-control devices, however, means immediate loss of control.

The possibility of short-circuit SCR failure is but one reason why the maximum dc output voltage from the system should correspond closely to the full-conduction angle of the primary SCRs. If a range of output voltages is required, it is a good idea to have a transformer tap-changing switch. The very real possibility of seriously overvoltageing the load could result from a design that operates normally with significant SCR phase-back.

As mentioned earlier, the SCR is not the only high-power half-control switch. Figure 15-4 shows the simplified schematic of a 2.25-MW, 150-kVdc primary phase-controlled power supply that first went into service with ignitrons as the control devices. (This power supply still exists as part of a high-power micro-

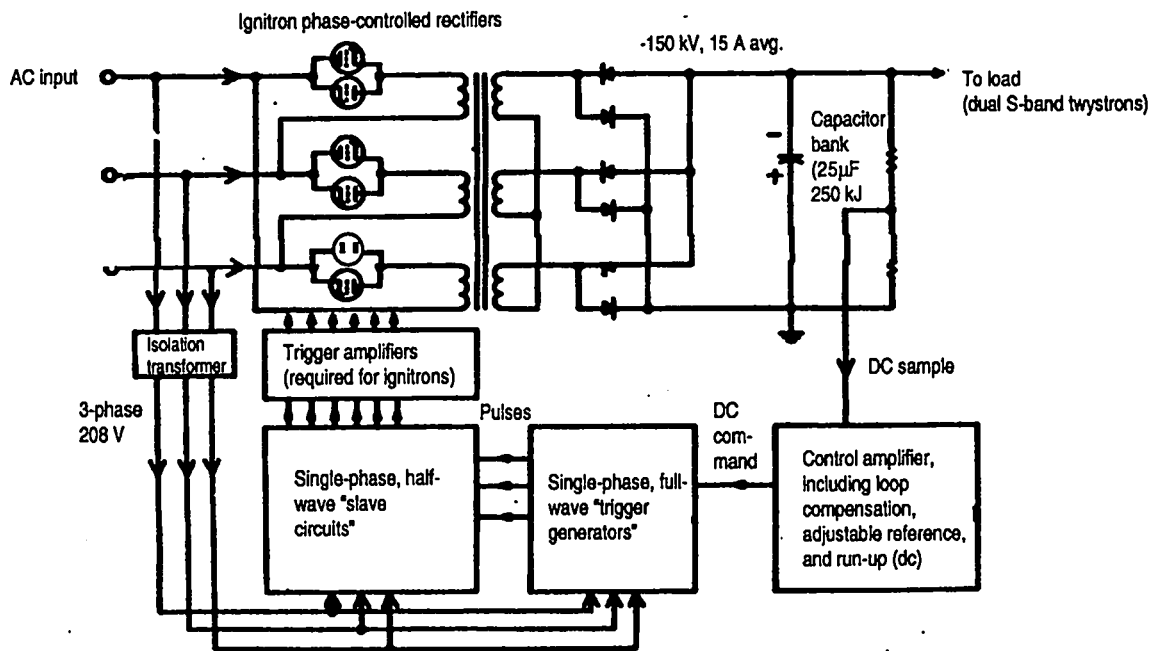


Figure 15-4. High-power variable dc power supply with primary phase control using ignitrons.

wave-tube test facility, but its original ignitrons have been replaced with SCRs.)

In the circuits of both Figs. 15-3 and 15-4, the half-control switches are shown connected within an "open-delta" transformer primary connection. This is used because the stress on individual switches is lower in this connection than if they were outside of the delta in series with each input line. They handle only the current in each delta leg, which is smaller by a factor of $\sqrt{3}$ than the line current. On the other hand, they must handle the full line voltage. If the delta winding is closed and the SCRs are in series in the line, two switches will be in series to block the line voltage at any given time. Having two switches in series, however, is less of a factor than it might appear because the sharing of reapplied voltage between the two may not be perfect.

Most primary phase controllers are outside of the delta of the rectifier transformer so that they can be part of an assembly that is completely separate from the transformer/rectifier. In fact, complete high-power, high-voltage dc systems often comprise three completely separate assemblies:

1. A full-power transformer that steps down whatever the subtransmission line voltage might be (4160 V, 13.2 kV, etc.) to the ~ 600 V that is compatible with single-device (per phase) controllers;
2. The primary phase controller itself; and
3. The high-voltage transformer/rectifier, which is often mounted in an oil-filled enclosure.

The waveforms of Fig. 15-5 show how the output voltage of the controller (transformer primary voltage) grows to resemble more and more the input line voltage as the conduction angle is increased. For small conduction angles, neither the output voltage nor the line current is especially sinusoidal in wave




Figure 15-5. Wave shapes of primary phase-control circuit as conduction angle is varied.

shape. The fundamental-frequency component of line current is also displaced in phase from the line voltage. Figure 15-6 shows the effect on the input-power factor as the output voltage is varied through phase-control alone. The figure illustrates why it is a good idea to have the useful range near the maximum




Figure 15-6. Effect on input-power factor as output voltage is varied.




Figure 15-7. Multiple-gang, three-phase (nine per phase) variable autotransformer, rated at 244 kVA.

output value.

Despite some very real shortcomings—including circuit complexity, line-current harmonic pollution, and the ever-present possibility of catastrophic failure—the SCR-type primary controller of high-voltage power supply is presently the most popular design by far. Almost all new requirements specify this type of controller. The popularity of designs based on the SCR is due partly to the electronic agility of SCRs and partly to its low cost and small physical size.

15.3 The variable autotransformer

At the other end of the technological spectrum, and sharing none of the advantages (or disadvantages) of the fully electronic voltage control methods discussed, is the archetype of all ac variable-voltage devices: the variable autotransformer. (It is known to many by the non-generic term "Variac," which is no more than a contraction of the terms "vary" and "ac.") In applications where the input volt-ampere demand is less than 100 kVA, the variable autotransformer, or suitable arrays of multiple autotransformers, can be quite economically competitive. Figure 15-7 shows just such a multiple-gang, three-phase array. All of its pickup-arms are driven by a single reversible motor by means of a chain drive. The outline drawing is of a 27-gang unit rated at 244 kVA. It has nine

transformers connected in parallel for each of the three phases. Because there can be no assurance that the tap voltage will be precisely the same for each unit, the outputs are collected through paralleling transformers, as shown in Fig. 15-8, that losslessly produce equalizing voltages so that there are no circulating currents.

Whether the variable autotransformer is technologically competitive with electronic control devices depends upon what is important for the particular application. If line-current harmonic distortion matters, this device produces none (or only a small amount due to non-sinusoidal magnetizing current of the core). If there is a fear that components downstream might suddenly be exposed to maximum-output voltage as a result of a catastrophic short-circuit failure in an active control device, this device won't do that. If the speed of response is not important, these devices can be incorporated in voltage- or current-control closed-loop regulators. (There have actually been electronic phase-control, high-voltage systems that have been retrofitted with autotransformers because of the problems associated with electronic control devices.)

15.4 The variable-voltage transformer

The variable-voltage transformer, or VVT, logically extends the circular-format variable autotransformer to higher current capabilities (up to 1 MVA). As shown in Fig. 15-9, the winding has been uncoiled and made flat. The advantages to this are obvious. When the winding is flat, its wires can be made flat, thus greatly increasing the pickup-brush surface area and even making dual-brush-pickup assemblies practical, as shown. If this is such a great idea, why haven't all variable autotransformers been made this way? The answer is leak-




Figure 15-8. Electrical schematic of multiple-gang variable autotransformer showing use of paralleling transformers to avoid circulating currents.





Figure 15-9. The variable-voltage transformer (VVT).

age inductance. The lack of tight magnetic coupling to the entire length of a linear winding is much greater with a flat-winding format than with the toroidally wound circular format.

Fortunately, this problem has been solved. (The solution has been patented, making this a unique device.) A conventional isolation transformer, but one with multiple secondary voltage taps, has been added to the linear winding. These tap points are connected to the corresponding voltage points along the linear

winding, in effect short-circuiting the leakage inductive reactance and greatly improving the voltage regulation over the entire range of output voltages. With dual brushes contacting adjacent conductors, there will be a small voltage difference between, of course. This voltage difference is equalized in the patented design by means of a balancing transformer, which rides along with the brushes. (This problem was handled differently in a competitor's linear-format variable transformer. It placed a pair of anti-parallel-connected silicon rectifiers in series with each brush and the common point. Because the contact voltage of the rectifiers is slightly greater than the turn-turn winding voltage, there was no circulating current.)

The motor control of these variable-voltage devices is no more complicated than that required for the circular-format transformers, and manually operated versions with hand cranks have been built for the absolute minimum of system complexity.

15.5 The step regulator

In the realm of very-large-power, variable-voltage devices is the under-load step regulator. Its output is not continuously variable, as was the case with all of the previous devices. It also differs in that its output is not a variable portion of the input voltage over the range from zero to 100%. Note from the schematic diagram shown in Fig. 15-10 that the typical step regulator comprises a trans-




Figure 15-10. The 32-step (16 buck and 16 boost) voltage regulator.

former with primary, or exciting, windings, which shunt the incoming line, and multi-tapped secondary windings, not unlike those of the VVT, called the series windings. Either one end or the other of each series winding is connected to the high side of each incoming line by means of a reversing switch (not shown in the schematic). When the winding is in one polarity, the voltage induced in it will add to, or "boost," the incoming line voltage. In the other polarity it will subtract from, or "buck," the incoming line voltage.

The taps on the windings are connected to contacts of a double-brush, make-before-break, tap-selector rotary switch (shown linearly in the schematic). One of the two brushes is always in contact with the source winding. At every other step, the two brushes will be connected across the entire voltage between tap points. A balancing, or "step," reactor produces an output that is the mid-point voltage. Typically there are 16 steps in each polarity, or 32 steps altogether. A common industrial application of the step regulator is as a line-voltage corrector that can buck or boost capability by 10% in discrete steps of $5/8$ of a percent. The regulator itself handles only 10% of the connected volt-amperes. A device rated at 2000 kVA (or 2 MVA), as illustrated in the figure, will regulate a 20-MVA system.

However, for the high-power transmitter service that we are interested in, full-range control is usually required—which means 100% buck or boost. Note that 100% buck will only be zero volts if the full series-winding voltage is precisely equal to the incoming line voltage, which it may not be. Also note that the maximum output voltage (100% boost) will be very nearly twice the incoming line voltage, not 100% of it, as was the case for the previous devices. (There are precious few things more embarrassing than designing a full-range control system with a buck/boost device to feed a rectifier transformer whose maximum-output voltage corresponds to the nominal line voltage. But it has happened. And more than once.) Even with full-range control, the step-regulator VA rating is only half the maximum-system VA rating because it only processes half of the maximum voltage.

With a standard 32-step device, there will be granularity in the stepped output of approximately 3% per step. For a dc power supply of 100-kVdc output, this amounts to about 3 kV/step. The rate at which the regulator will slew from minimum to maximum output is user-selectable, with 30-s run-up time being fairly typical. In systems where automatic return-to-zero is implemented, the run-down time can be made more rapid, such as 15 s. If such granularity is too great or the speed not great enough for a particular application, the step regulator is not the device to use. Nevertheless, the step regulator is the least expensive means of obtaining voltage variability at very high VA ratings. (Although the word "is" should probably be replaced with "was" in the last sentence; it is not known if there is a domestic source for such devices at this time.)

15.6 The induction voltage regulator

The induction voltage regulator, or IVR, is often referred to by the non-generic name of "Inductrol." It is also a buck/boost device, but it has continuously variable output voltage. If the step regulator were a Chevy, the IVR would be a Cadillac. Like the step regulator, it has shunt and series windings, as illustrated




Figure 15-11. The induction voltage regulator (IVR).

in the schematic diagram of Fig. 15-11. The shunt windings, however, are part of a rotor assembly that can be rotated with respect to fixed, or stator, series windings. One end of each series winding is permanently connected to the high-side of each incoming line.

Depending upon the angular orientation of the rotor with respect to the stator, the voltage induced in the series windings will vary in amplitude and polarity. One complete revolution of the rotor will generate every condition from maximum buck to maximum boost. The turns ratio of the shunt and series windings will determine the degree of output buck and boost, from 100% buck/boost for full-range control, down to 10% buck/boost for line-voltage correction or filament-voltage regulation, as indicated in Table 15-1. Still, obtaining true zero-voltage minimum is even trickier with the IVR than with the step-regulator. Not only must the series-winding voltage exactly match the incoming-line voltage, but the rotor must be stopped at exactly the proper spot. Mechanical stops must be perfectly positioned and then never be budged—neither of which is entirely likely. (Electronic return-to-zero circuits have been built with non-linear dynamics that greatly exaggerate IVR output voltage as zero is approached. This method does as well as any in achieving minimum, if not zero, voltage.)

Both the IVR and the step regulator can be incorporated in closed-loop voltage- or current-control servo systems. Both require “bandwidth” compromises to avoid continuous “hunting.” (Bandwidth here refers to a voltage band rather

Table 15-1. Characteristics of the induction voltage regulator.

Device buck/boost ratio	Ratio of minimum and maximum output voltages to input	Ratio of load VA to device VA
10%	0.9/1.1	10
20%	0.8/1.2	5
50%	0.5/1.5	3
100%	0.0/2.0	2

than a frequency-response band.) Experience with large IVRs has shown that the rotor setting is not completely insensitive to mechanical vibration. (When operating in its "killer" 50%-duty mode at a 10 pps rate, the Haystack Hill LRIR transmitter, which has been referred to several times already, produced sufficient floor vibration to cause very gradual but very noticeable walk-off of the output of the IVR that fed the high-voltage dc power supply.)

In external physical appearance, there is little difference between the IVR and the step regulator. At the highest VA levels, both devices are enclosed in oil-filled tanks for electrical insulation and cooling. Some IVRs, however, have been built for lower VA service, where they are competitive with variable autotransformers and VVTs, even though they are most costly. At these lower power levels, they are air-insulated and convection-cooled.

Like the step regulator, the IVR, or at least the type that goes by the trade name "Inductrol," is no longer available domestically. However, a form of it is still being manufactured abroad.

16. Electronic Voltage Regulators

Some microwave tubes require more precise regulation of beam voltage than can be obtained from the power supplies and regulators already discussed. This is especially true if the nature of the transmitter output waveform is continuously changing, which is often the case with modern sophisticated radar systems. High-frequency, switch-mode power supplies, which will be described later, can outperform the most precise variable-voltage device, which uses conduction-angle controlled rectifiers simply because of their usually higher self-generated operating frequency.

No regulating system can outperform the electronic variable-resistance voltage regulator, either in precision or in transient response. In the case where intrapulse voltage regulation is required, which is usual for a long-pulse system where the intrapulse voltage droop obtainable from a capacitor bank of practical size is excessive, there is virtually no alternative to the electronic variable-resistance regulator.

The variable resistance in such regulators, which are also sometimes called "linear" regulators, is supplied by a "series," or "pass," device that is connected in series with the high-voltage loop. This device must be "full-control," such as the ones previously discussed in regard to hard-tube pulse modulators. In large-power, high-voltage systems, the series device is customarily a large-power vacuum tube. The tetrode is often chosen for this application because adequate anode, or load, current can be obtained with a grid-cathode voltage that never needs to be positive. This situation greatly simplifies the design of the error amplifier that provides the beam voltage. (The power MOSFET has characteristics that are nearly ideal for such service, but its use is limited to low- or medium-power situations because it is considerably more difficult to harden against fault conditions in the high-voltage loop.)

16.1 Voltage regulators that handle system average current

Figure 16-1 shows a typical medium-power radar transmitter's high-voltage system that uses a variable-resistance electronic regulator. The peak-pulse current is supplied from a capacitor bank so large that intrapulse regulation is not required. The electronically controllable variable resistance in series with the high-voltage loop of the microwave tube comprises the parallel-connected tetrode vacuum tubes *V1-V3* and is connected between the high-voltage power supply and the capacitor bank. The microwave tube used in this system is a 35-GHz coupled-cavity TWT that produces 30-kW peak-power output. The optimum beam voltage is between 47 and 48 kVdc, but it must be precisely controlled and free from ripple-frequency components because of the high phase-pushing factor for this type of high-gain, broad-band tube. In addition, system specifications required that phase-modulation sidebands be -50 dBc.

To give some idea of what this means in terms of regulator performance, consider that for each phase-modulation component to be 50 dB below the carrier, its peak voltage must be 0.003 of the carrier. Assuming that the ripple lines

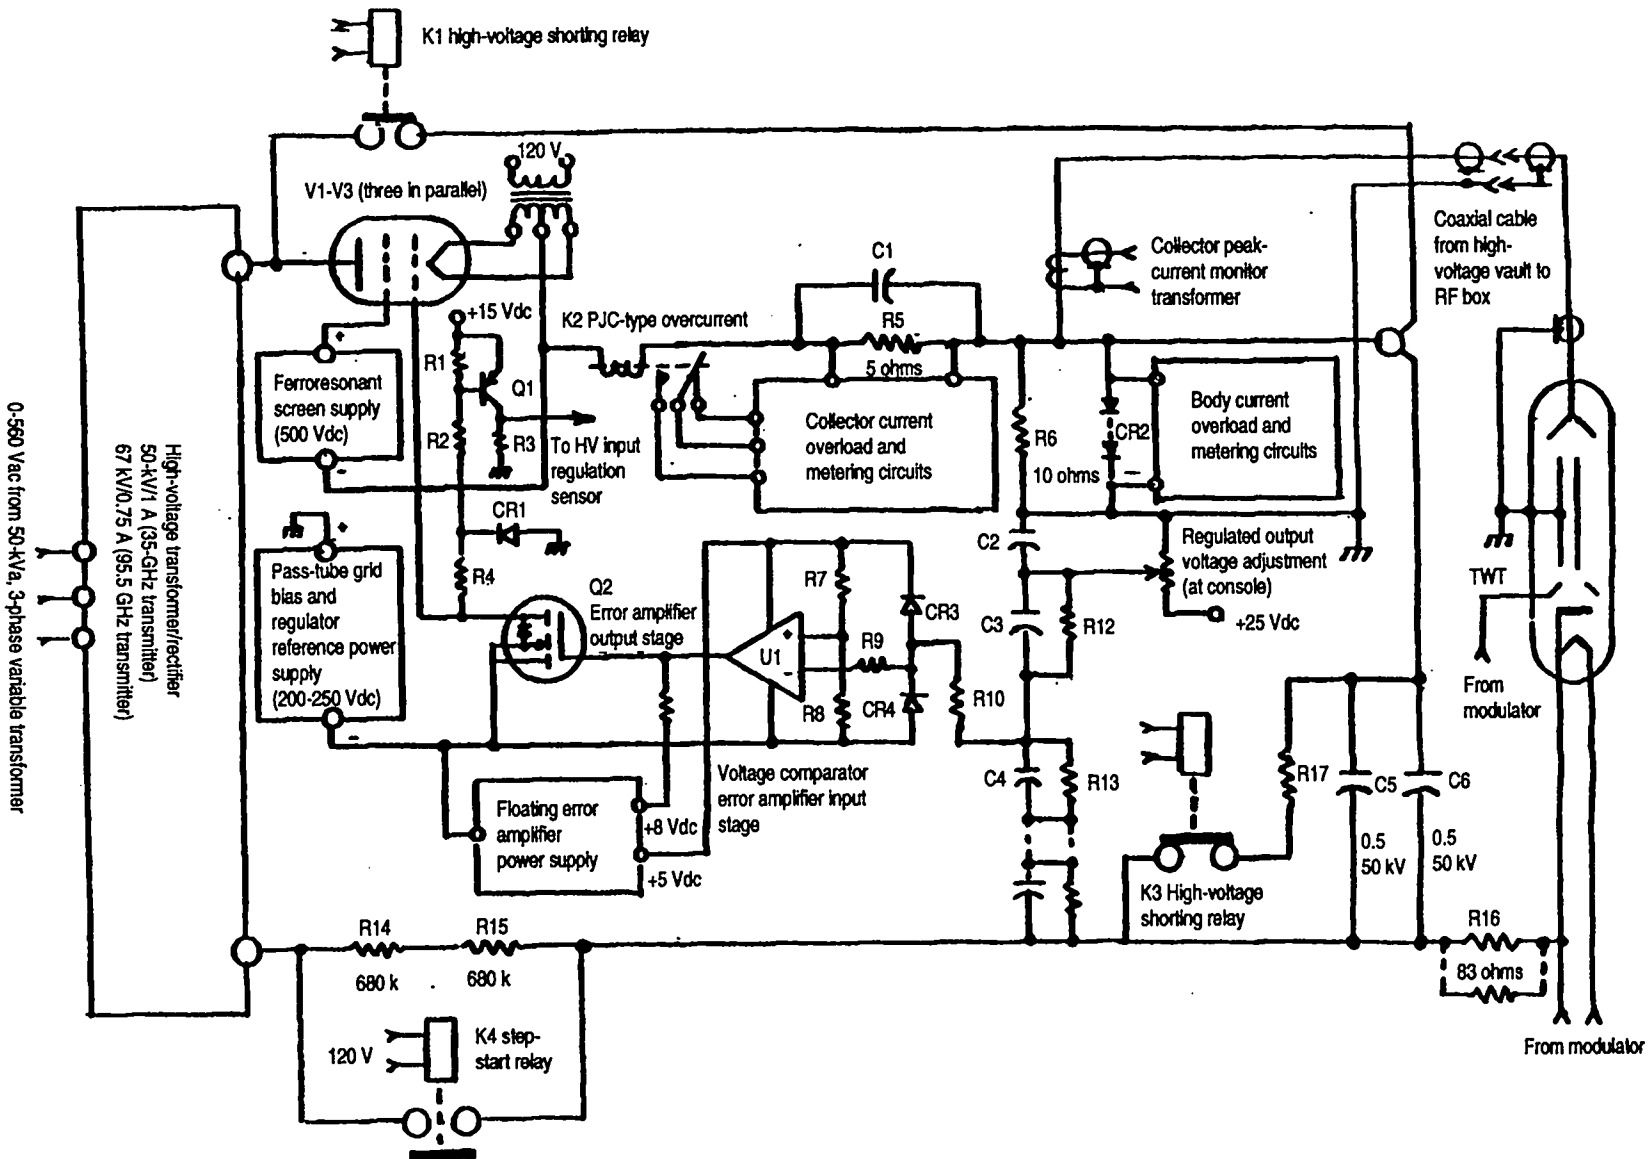


Figure 16-1. Example of electronic voltage regulator that handles microwave-tube average current.

are discrete and harmonically related to the 60-Hz primary-power frequency, we can expect there to be paired sidebands, with each sideband of a pair being above and below the carrier by the same frequency difference. The voltages of the two sidebands will add in quadrature with the carrier voltage so that the sine or tangent of the peak angular modulation will be 2×0.003 , or 0.006. The corresponding phase angle is 0.36° . The phase-pushing factor of a TWT of this type is typically 50%/percent of beam-voltage change. The allowable peak ripple voltage, therefore, is 0.36° divided by 50%/percent, or 0.007% peak ripple voltage, which is not much.

Note that the anode-cathode circuit of the regulator series tubes is in the positive return of the high-voltage power supply. The cathode of the regulator series tube is connected to the collector lead of the TWT, through which passes the great majority of the TWT cathode current. The collector lead is connected to ground through resistor $R6$. The body of the TWT is also connected to ground. So are its waveguide input and output connections, which are an integral part of the TWT body structure. The voltage across $R6$ is proportional to the current intercepted by the TWT body, which is normally a small fraction of the total beam current. However, in the case of an internal TWT gun arc, all of the current is body current, which is limited only by $R16$ (83 ohms) to a value of 50 kV/83 ohms, or 600 A. The series-connected diodes, shown as $CR2$, shunt the fault current around $R6$ while maintaining a total forward diode-drop great enough not to distort the voltage corresponding to normal body current.

The upshot of all of this is that the cathode of the regulator tube is very near ground potential, which is also the common point for the sample of the TWT cathode voltage obtained from the compensated voltage divider comprising $R13/C4$ and $R12/C3$. Of course, the criterion for frequency compensation is that the product of $R13$ and $C4$ equals the product of $R12$ and $C3$. The error amplifier topology can be relatively straightforward because there is no common-mode voltage difference to overcome. In this design, however, there is at least one unique feature. The common terminal of the error amplifier has been deliberately offset from ground by an amount equal to the negative voltage required by the error-amplifier output stage to assure current cutoff in the regulator series tube, which in this case is a minimum of -200 Vdc. This power supply, therefore, is a precision low-ripple dc source.

By offsetting the error-amplifier common terminal, the attenuation factor of the sample-voltage divider can be made many times smaller than with a grounded error amplifier. This offset is less important for the dc performance than it is for ripple rejection because the sampled ripple components are many times greater. The attenuation factor in this case is 50,000 V/200 V, or 250, whereas in a more conventional regulator the attenuation factor might be 50,000 V/5 V, or 10,000, if a 5-V reference is used for the error amplifier. The actual reference voltage used in this design is 2.5 V, but it floats atop the 200-V grid bias supply. The input stage of the error amplifier is a voltage comparator, $U1$, whose signal gate is protected by the resistor/diode network, $R8$, $R10$, $CR3$, and $CR4$. The error-amplifier output stage is a high-voltage MOSFET, $Q2$, which produces variable grid voltage to the $V1-V3$ tetrode by varying the current through $R4$ in response to the difference in voltage between the positive and negative gates of $U1$. The tetrode grid

variable voltage produces the required variable resistance in the high-voltage loop to precisely regulate the TWT cathode voltage.

The TWT beam current is in the form of 6-A peak-current pulses that do not exceed 50 μ s in duration at duty factors up to 10%. The pulse current is obtained from a relatively small capacitor bank, C5 and C6, totaling 1 μ F. The intrapulse voltage droop is acceptably small. The voltage regulator, therefore, needs to pass only the maximum average beam current, 0.6 A, while maintaining constant a pulse-to-pulse voltage across the capacitor bank.

Like many high-voltage dc conditioning systems that have electronic regulators, this one has an ac variable-voltage device as well, a motor-driven variable autotransformer, that permits the dc output from the transformer/rectifier to be gradually applied. Or it can also be reduced to a lower value (including zero) for diagnostic purposes. In recovering from a high-voltage shutdown resulting from an internal TWT arc, the power supply can be "slapped on" with full ac input. The peak capacitor charging current and the rate-of-rise of TWT cathode voltage is determined by the "step-start" resistance, R14 and R16. When the operating voltage is reached, this resistance is short-circuited by the high-voltage relay K4, and the operation of the electronic regulator is sensed by the flow of current through R4, which is in series with R1 and R2 as well. This current turns on transistor Q1, signaling the control circuits that the regulator is functioning.

However, before the regulator begins to function at dc input voltages lower than the normal operating value, the tetrode series tubes are zero-biased. There is also no anode-cathode voltage, which means that the screen grid will attempt to play the role of surrogate anode in the collection of current from the cathode. Screen current must be externally limited to prevent damage to the screen grids. A most effective way of doing this is to operate the screen supply from the output of a ferroresonant constant-voltage transformer, which is matched as closely as possible in VA rating to the normal full-load requirements of the screen supply. When such a transformer is short-circuited, its output current is typically only one-and-a-half to two times as great as the normal full-load current, and perhaps surprisingly it runs cooler under this condition than under full load. (The worst case is when it is open-circuited.)

A less expensive way of limiting screen current is to use ac capacitance in series with either the screen-supply ac input or ahead of the rectifier. The effect on screen-voltage regulation of capacitive reactance that is capable of limiting short-circuit current to two times full load can be quite modest. The ferroresonant transformer has the obvious technical advantage that it not only limits short-circuit current but it regulates source voltage in the normal operating range.

16.2 Voltage regulators that handle system peak current

When the transmitter waveform dynamics are such that an energy-storage capacitor bank would be impractically large in order to limit intrapulse voltage droop to acceptable levels, there is virtually no alternative to the pulse-current voltage regulator. The Haystack Hill LRIR is a classic example of such a transmitter (see Fig. 16-2). Its four parallel TWT high-power amplifiers draw a total intrapulse current of 40 A for pulse durations up to 50 ms. This is a per-pulse charge transfer of 2 C, which is large enough to completely drain capacitor banks

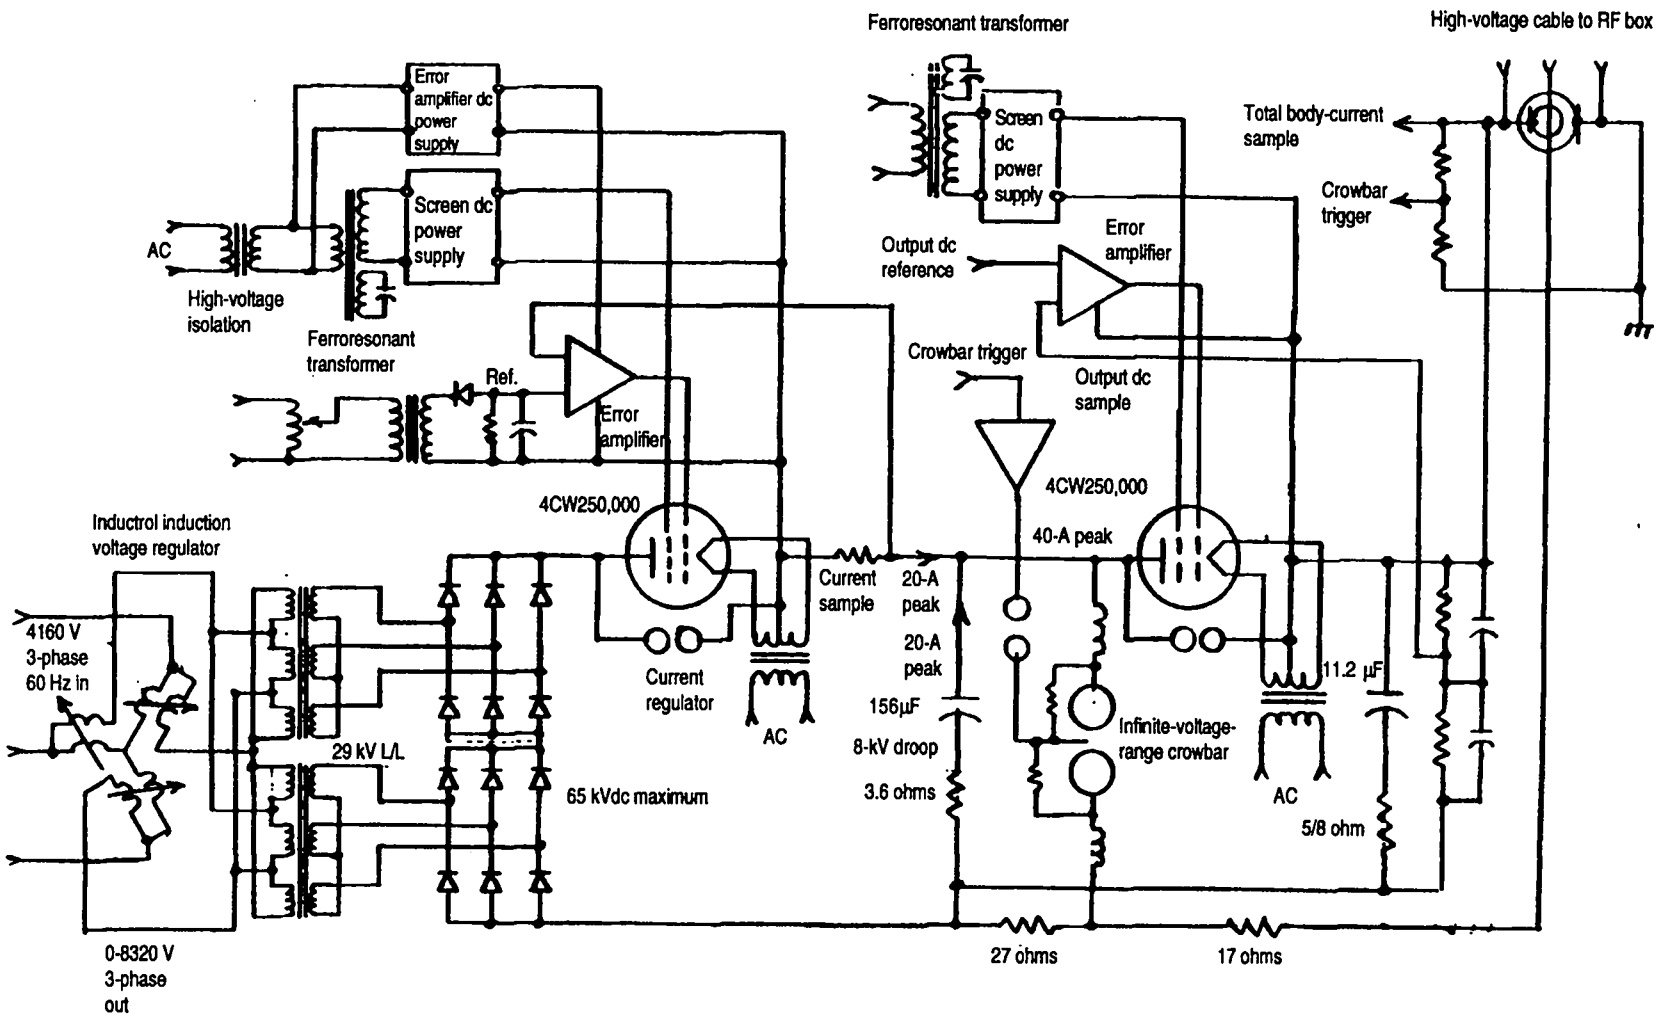


Figure 16-2. Example of high-voltage system having an electronic voltage regulator handling microwave-tube peak-pulse current and an electronic regulator handling average current.

of common size. Even though the LRIR transmitter main capacitor bank of 156 μF could hardly be considered commonplace, the removal of 2 C from it would cause its voltage to droop by almost 13 kV. However, the dc power supply is capable of putting out an average current of 20 A, which it did whenever the transmitter was operated at its maximum duty factor of 50%. Therefore, the entire 40-A peak-pulse current did not have to come from the capacitor bank, because half of it would come from the dc power supply. Even so, the capacitor-bank droop was over 6 kV per pulse.

A voltage regulator was provided that would remove the 6-kV voltage change while passing the full 40-A peak-beam current. The output voltage applied to the cathodes of the TWTs remained within approximately 30 V of the desired operating voltage, approximately 47 kV. (Tetrodes are inherently good source-voltage regulators because of their high incremental anode resistance or constant-current nature.) The anode voltage of the 4CW250,000 tetrode series regulator tube started each pulse 6 kV higher than at the end. The total source voltage at the beginning of each pulse was 47 kV + 6 kV + 3 kV for a total of 56 kV. This value takes into consideration the headroom required for the tube to pass 40-A peak without excessive screen-grid current.

Apart from the obvious requirement that the series-connected device in a pulse-current regulator must be capable of passing the full-peak-pulse current of the microwave tube(s)—a conventional voltage regulator must merely pass the time-averaged value of the total beam current—the major difference between the two types of regulators is transient response. Whereas the series tube of a conventional regulator in the linear mode is in a state of almost-continuous conduction—the degree of which changes only slowly in response to changes in transmitter duty factor—the conduction of a pulse-current regulator must change from full-off to full-on during each pulse. In addition, it is precisely the intrapulse grid-cathode voltage of the regulator tube that keeps the microwave-tube beam voltage constant. The error amplifier, therefore, must be both precise and swift.

Fortunately, the timing of the load-current pulses is not unknown to the designer. So, often the slew rate of the error amplifier is assisted by means of an open-loop synchronous switch that receives its timing information from the same source as the pulse modulator. Other than that, there is no reason why the error-amplifier topology of the pulse-current regulator would be significantly different than that shown in Fig. 16-1. In that figure, the error-amplifier output resistor R_4 and the current-handling capability of the output stage, Q_2 , must be sized so that the input capacitance of the regulator tube will be charged and discharged so rapidly that the buffer capacitor bank, which is at the output of the voltage regulator, is not excessively discharged before the regulator tube can take over the delivery of the full-load current. Needless to say, due to the higher gain and speed and lower internal impedance of the pulse-current regulator, this design is more difficult to stabilize against oscillatory instability than a conventional regulator. It is also essential to provide for high-frequency bypassing of the filament leads of a directly heated cathode tube, which most high-power tetrodes tend to be. (Although this feature is not shown in Fig. 16-2.)

In addition to the electronic voltage regulator in the LRIR transmitter, some means was also required to keep the dc power supply from being virtually short-

circuited by the linearly discharging capacitor bank. In which case it would attempt to supply the entire 40-A peak current by itself. The solution, as shown, was another electronic current regulator set to pass no more than 20 A, which it did continuously at a 50%-duty factor. It required anode-voltage headroom of approximately 2 kV to pass 20 A, so the total dc-power-supply output voltage was 58 kV. The anode drop of the current regulator increased throughout each pulse as the voltage-regulator anode-drop and capacitor-bank voltage both decreased.

Unlike the voltage regulator, the cathode of the current regulator was not at virtual ground potential. All of the circuits referenced to its cathode had to be isolated for at least 10 kV and had to tolerate the same rate-of-change of voltage as the main capacitor bank, which was $I/C = 20 \text{ A}/156 \mu\text{F}$, or approximately 105 V/s. The error-amplifier inputs, however, required no elaborate level-shifting or common-mode rejection, because circuit current could be sampled by means of a series resistance wherever it was most convenient. As shown, this point was in series with the cathode of the current-regulator tube.

Not all systems that require pulse-current voltage regulators also need current regulators. If capacitor-bank droop can be maintained at 1-2% of TWT beam voltage, a pulse-current regulator will not be required. If the actual capacitor-bank droop is only 1-2 kV or so, it is unlikely that the dc power supply will hog the pulse current. A dc inductor in series with the power supply can be made large enough to keep its output current constant. In so doing, however, the inductor can also store a great deal of energy that must be dealt with one way or another during fault conditions. Often this is a difficult problem.

16.3 Voltage regulators that do not handle the full microwave-tube beam current

The critical voltage for all microwave tubes is the voltage between cathode and ground because the RF circuit for all practical tubes is grounded. The only microwave-tube current component that must flow in ground is that which is intercepted by the body because the electron beam has been imperfectly focused. If the collector of the microwave tube is isolated or insulated from ground, then the collector current can be conducted in a path that is separate from the body-current path. (Microwave tubes, of course, can have multiple collectors, in which case there will be multiple collector-current paths.) Therefore, the series tube of a voltage regulator that is regulating the voltage between cathode and ground needs to handle only the current intercepted by the body of the microwave tube, which in a well-designed tube will be less than 10% of the total beam current.

A good example of such a regulator is shown in Fig. 16-3, which is a simplified schematic diagram of the high-voltage circuits for high-power RF amplifier groups. Each group has eight high-power (200-kW peak power) TWTs operating in parallel with respect to the high-voltage system. The nominal peak-pulse cathode current is 16 A for each TWT, 128 A for all of them. A pulse-current linear regulator for this much current would be awesome, indeed. So would any capacitor bank used to replace it because its maximum pulse duration is 2000 μs .

However, the TWT—which is specifically designed for long life, a high degree of phase linearity, and high efficiency—does not have one isolated collector but

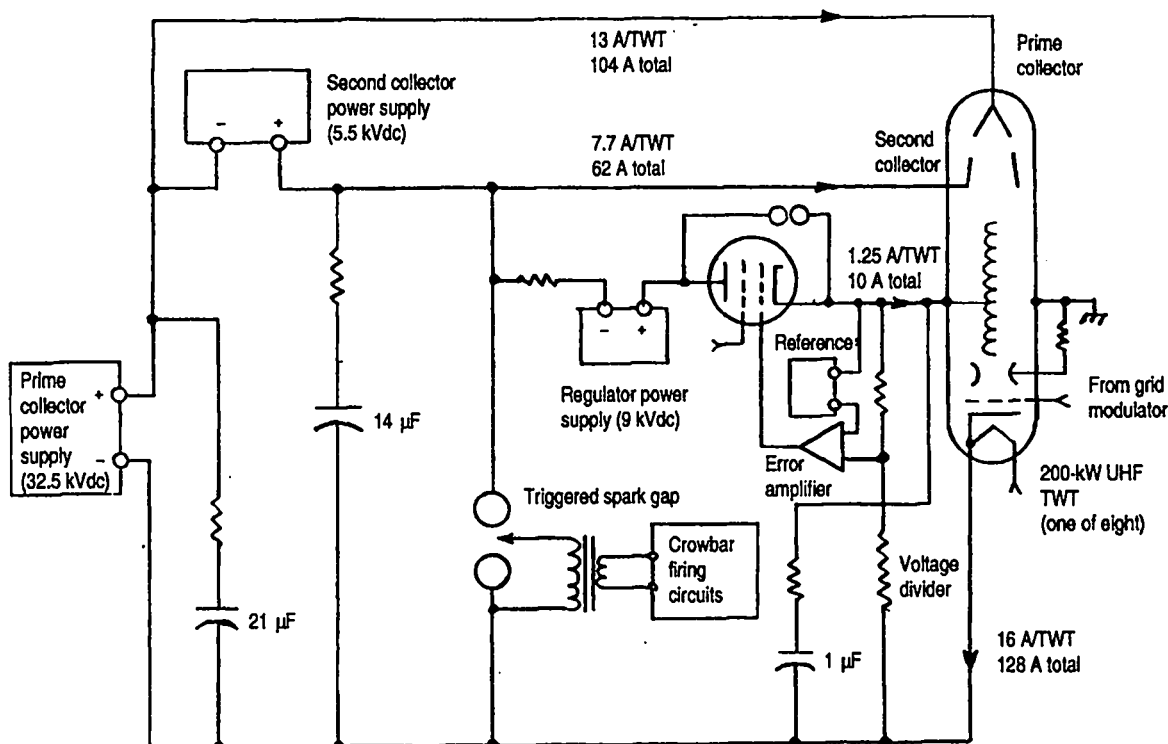


Figure 16-3. Electronic linear voltage regulator that handles only TWT body current.

two, both of which are depressed in voltage with respect to the grounded helix-type circuit. Three dc power supplies are required: one at 32 kVdc, which is connected between the cathode and prime collector; a second at 5.5 kVdc, which is connected between the prime and second collectors; and a third at 9 kVdc, which is connected between the second collector and the anode of the tetrode series tube of the linear voltage regulator. The resistance of the voltage regulator is electronically varied to maintain the voltage between cathode and ground of the TWT at its nominal value of 41 kVdc. However, the current that passes through the regulator tube is only about 8% of the total peak-cathode current, or 10 A. A modest-sized tetrode (4CX5000 class) is adequate for this purpose.

16.4 Regulators that must also be modulators

The regulator circuits discussed so far are used in transmitters having separate pulse modulators, either modulating-anode or control-grid pulsers. But such regulators are also used in transmitters that have microwave tubes with diode-type guns. In order to pulse-modulate them, either a line-type pulser or a hard-tube switch is required that is capable of handling the full beam current and voltage of the diode electron gun. If active pulse-top regulation is a requirement, the line-type pulser is not an option. However, we have already discussed voltage regulators that must handle the full peak-pulse current of the microwave tube. If, in addition, the series regulator tube is capable of holding-off the full diode beam voltage as well, the roles of pulse modulation and pulse-top regulation can be integrated into the same circuit.

Figure 16-4 shows a simplified schematic of a practical pulse-current voltage regulator in which the regulator tube is in series with the positive return bus of

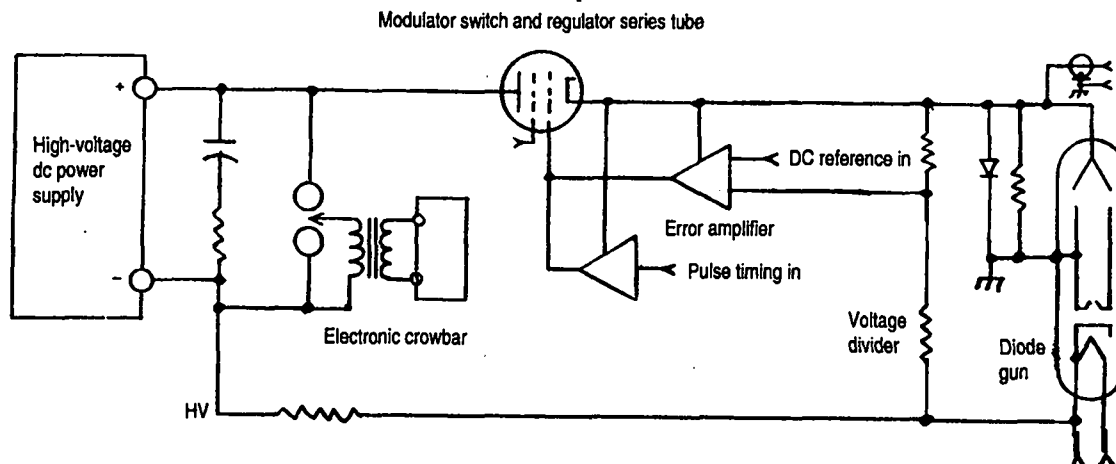


Figure 16-4. Modulator switch and regulator series tube in positive return of high-voltage power supply.

the high-voltage circuit. The advantage of this location in the circuit is that the common nodes for the error amplifier and the voltage sample are both virtually at ground potential, so that direct-coupled, high-speed circuits can be used throughout. Presumably, if the voltage hold-off capability of the regulator tube is adequate, it can be pulsed or biased to the non-conducting state during the interpulse interval by means of an amplifier channel that overrides the error-amplifier channel in response to low-level pulse-timing information, as shown. The combination modulator switch and regulator tube is shown as a tetrode because it is highly advantageous that the control grid operate in the negative-grid regime under all conditions.

Although such a connection can presumably be made to work, it has one usually crucial disadvantage. When the switch tube is off, the full power-supply voltage appears across it and not across the diode electron gun. Thus, the negative terminal of the power supply (and capacitor bank) is at ground potential, and the positive terminal is above ground by the full system voltage. When the switch tube is turned back on, the voltage across it drops to the minimum head-room requirement, pulling the positive terminal of the high-voltage system back to near ground and the negative terminal down to the operating cathode voltage of the diode gun. This is very stressful for the high-voltage power supply, the capacitor bank, and even the electronic-crowbar circuits. It should be avoided.

Figure 16-5 shows the combination of modulator switch and regulator tube in the circuit location most common for a hard-tube modulator switch: in series with cathode of the microwave-tube's diode electron gun. In this location there is no common terminal shared by the beam voltage sample, the error amplifier, and the grid-cathode circuit of the regulator tube. In fact, the error amplifier and voltage-sample common are separated from the regulator-tube grid-cathode circuit by the entire high-voltage power supply output. Modern high-speed data links using fiber-optic coupling are capable of bridging this voltage gap, as shown. Voltage telemetering is commonly performed by complementary voltage-to-frequency and frequency-to-voltage converters. Usually the variable-frequency pulse train is optically coupled.

However, the disadvantage of this topology is that the proportional-voltage

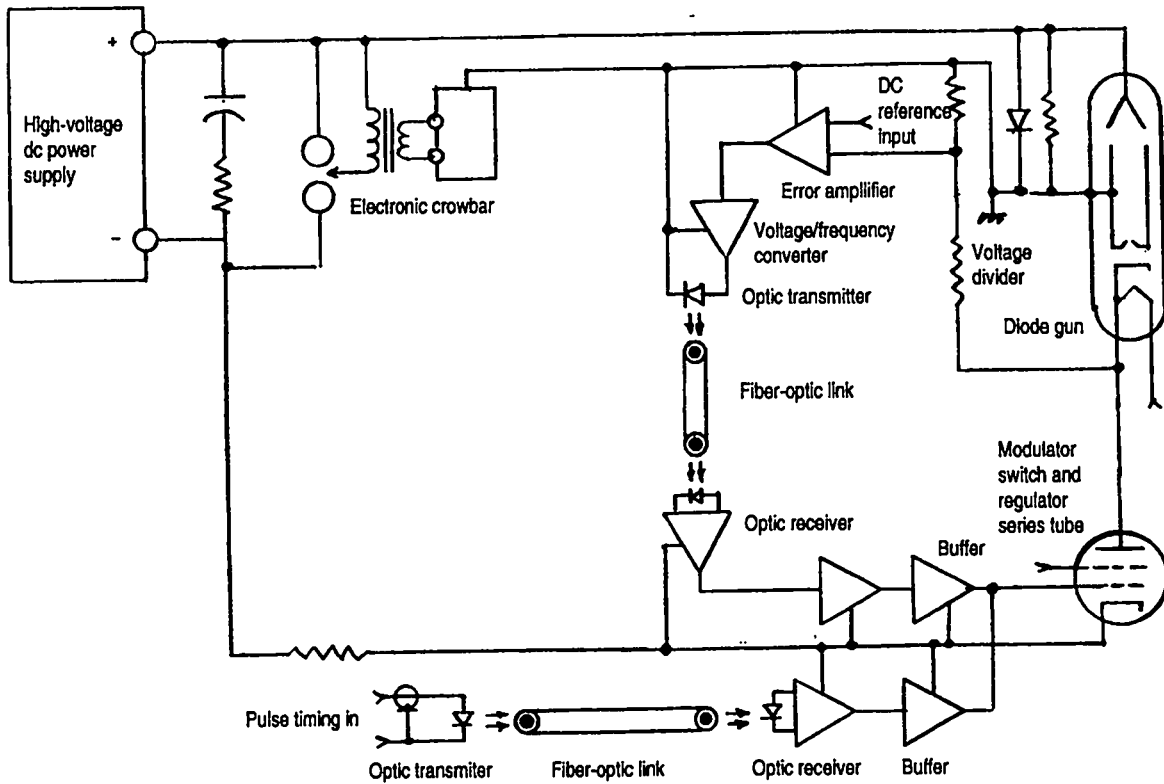


Figure 16-5. Modulator switch and pulse-voltage-regulator tube in series with high-voltage bus.

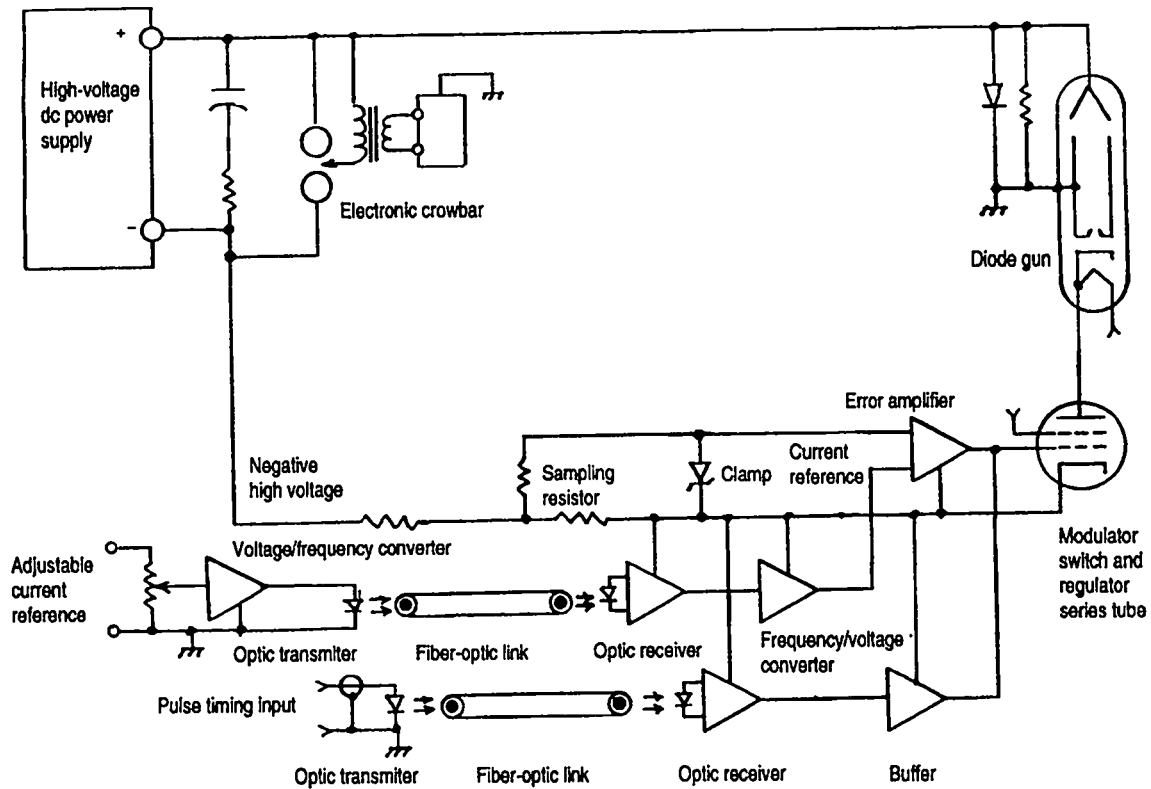


Figure 16-6. Modulator switch and pulse-current-regulator tube in series with high-voltage bus.

data link, no matter how high its speed, lies within the closed loop of the voltage-regulating servo system, thus limiting the speed of its response. To get around this problem, it is not unusual to use a proportional light-link whose optic-receiver output voltage is at least monotonic with optic-transmitter input current. (Absolute linearity is not a requirement. Monotonicity is.)

A permissible question at this point is, why try to regulate beam voltage in the first place? It may be better to attempt to regulate pulse-beam current instead, as shown in Fig. 16-6.

In a properly functioning diode electron gun, beam current is always related to the $3/2$ power of beam voltage. Regulating beam current to a given degree of precision will maintain beam voltage to a value that is one-and-a-half times as precise. Moreover, the beam current can be sampled by means of a series resistor, as shown. This sampling resistor restores the desirable common-terminal situation of the circuit in Fig. 16-4, only at the high side of the power supply rather than at ground. Both the pulse-timing information and a voltage proportional to the desired peak-beam current can be optically coupled to the high-voltage common node without impinging on regulator-response speed or open-loop gain. During the rise time of each pulse, the error amplifier will be saturated, permitting the switch tube to deliver all of the capacitance-charging output current at its disposal. As beam current approaches the desired value, the voltage across the sampling resistor will approach the current-reference voltage, and the regulator will attempt to maintain that condition until the pulse-timing information once again overrides the error-amplifier output and shuts off the regulator/modulator switch.

17. Switchgear and Substations

Switchgear and substations are not always matters of concern for transmitter designers, because they are often part of the facilities of a typical installation. However, this does not mean that a transmitter designer—or specifier, for that matter—can be ignorant of the characteristics of either. The substation, or what constitutes it, is often where the switchgear physically resides.

17.1 Switchgear

In general, switchgear comprise the devices that are necessary to make and break connections between the sources of primary power and the various loads that make up a transmitter system. These devices include power contactors, circuit breakers, and manual disconnects.

Usually electro-mechanical power contactors are normally open. Their contacts close in response to control power being applied to a magnetic solenoid. They open because of the force of gravity, a compressed spring, or both when control power is removed and the magnetic field of the solenoid collapses. Contactors of this type are usually slow to open (especially if they are air-insulated), often taking 5-10 cycles at a 60-Hz rate to completely clear the load circuit. As has been frequently mentioned in earlier sections, periodic fault conditions stemming from high-voltage breakdown will effectively short-circuit the feed lines to the high-voltage power supply of the transmitter. No matter how rapidly a circuit is mechanically disconnected, however, current will not cease flowing until it has been commutated off. Alternating current periodically passes through zero, which is a natural commutation point. Until it passes through zero, the conduction path will be maintained through a mechanically opened circuit by means of an arc between the contacts. Even if contacts could be made to open instantaneously, current could continue to flow. It would flow for another half-cycle for a single-phase or delta-connected three-phase rectifier feed, or for one-third cycle for a wye-connected three-phase rectifier feed.

The practical mechanical goal for a contactor's opening speed is to make it less than one-half cycle, which is about 8 ms for 60-Hz primary power. The switch mechanisms capable of doing this use vacuum-insulated contacts and are usually referred to as vacuum contactors. The vacuum insulation, with its greater dielectric strength, permits more closely spaced contacts for a given voltage-hold-off rating, thus lowering the opening time. A typical vacuum-insulated switch assembly is shown in Fig. 17-1. The fastest such mechanisms are either normally closed and therefore powered open, or they are normally open, in which case they have a high-speed mechanically coupled booster solenoid to assist the normal spring-driven opening. The opening coils in the highest-speed contactors are nominally low-voltage solenoids that are momentarily overvoltaged, usually by the discharge of an energy-storage capacitor bank through an electronic switch like an SCR. The triggering gate-signal for the electronic switch is generated by electronic monitoring circuits, which sense a microwave-tube fault-current overload and simultaneously fire the electronic crowbar, if one is used.




Figure 17-1. Section view of typical vacuum-insulated switch.

Of course, the normally closed form of contactor is not inherently safe, because it requires coil power to stay in the open state. Therefore, it must be preceded by a "nominally fail-safe" normally open contactor. This contactor should not be vacuum-insulated but air-insulated, because even a normally open, vacuum-insulated switch is not safe if its vacuum bottle has taken on air. (The term "nominally" is used to modify "fail-safe" in this case because at least one power contactor so designated became locked in the closed position when its contacts welded together. The contactor was finally opened by the less-than-instantaneous application of a crowbar—and not the electronic type.) The normally open type that has a high-speed auxiliary opening coil provides extended-life operation of the vacuum-contact bottles because, under normal shutdown conditions, the opening of the contacts is relatively slow thanks to the compressed-spring actuation. The capacitor-discharge, high-speed opening occurs only in response to transmitter high-voltage faults (which, we would hope, are not responsible for the majority of shutdowns).

A contactor of this type has two distinct current ratings or VA ratings. One is the amount of current that can be continuously conducted. This rating is related to long-term temperature rise of the contacts due to ohmic heating. The other rating is the amount of short-term fault current that can be terminated following the mechanical opening of the contacts. If fault current exceeds this rating, the

contactor cannot be expected to clear the circuit, because the arc between the open contacts may not be quenched when current passes through the zero point.

When primary-power input voltages are less than 600 V, solid-state relays are practical. If they are operated within their short-term (sub-cycle) peak current and heating-impulse ratings, they will clear the load circuit at the first current-zero following interruption of gate drive, which can be accomplished with electronic speed. Of course, if the power supply already has an SCR primary-voltage controller, the solid-state interrupter feature is included at no extra cost. Even if the power supply has a less-sophisticated means of voltage variation and the input voltage less than 600 V, a solid-state relay (SSR) should at least be considered. The SSR, however, should be treated with the same degree of caution as the vacuum relay, and it should always be backed up by a positive-action, air-insulated disconnect switch. (SCRs have been known to fail in this service, and when they do, their electrical properties bear a strong resemblance to the bus bars that connect them.)

Contactors, including SSRs, are closed by the application of control power. They will stay that way until the control power is removed. On the other hand, circuit breakers, which are usually located on the line side of contactors, may either be manually closed or closed by the action of a motor drive. Often in the process of closing, the motor drive charges powerful springs that store the energy required to open the breaker. Also unlike contactors, breakers will automatically open in response to a current overload, either because of thermal or magnetic effects. In addition, they can be used in a quasi-contactor mode when they are equipped with "shunt-trip" coils, which trip the breaker in response to the application of an external control signal. An even more useful feature is the "undervoltage trip." A breaker equipped with this feature will trip free when a control voltage, either ac or low-voltage dc, that is applied to the undervoltage-trip circuit falls below a specified minimum value. Such a feature can be incorporated in a "fail-safe" electrical interlock system, especially for personnel safety.

Circuit breakers have characteristic timing curves that relate the degree of current overload to circuit-clearing time. When these curves are properly coordinated with the circuit-clearing time of the high-speed contactor, the circuit breaker(s) will not open under normal fault-reaction conditions before the contactor does, but they will still protect major components and wiring even if the contactor completely malfunctions.

The other type of switchgear encountered is the manual load-break disconnect switch. Typically, this is a spring-loaded switch that resembles a knife-switch. It is operated by a polelike handle that is vertical and flush with the equipment front panel when the switch is closed. When it is pulled down in a 90° arc to the horizontal position, it is open, and it can be locked in that position. Most of the handle travel is used to charge the opening and closing springs, which then open and close the contacts quite abruptly near the endpoints of the operating-rod travel. The manual disconnect can be expected to have large fuses in series with it. These fuses, when properly coordinated with the switchgear, are the circuit-clearing devices of last resort. The combination is referred to as a manual-fused disconnect.

In systems without manual disconnects, draw-out circuit breakers can serve

the same function. Such breakers are “racked-out” of their equipment cabinets on slides or rails, physically disconnecting their terminals from the line and load connections within the equipment enclosure. Although more cumbersome and slower than the manual disconnect, they can provide the same, or better, ultimate load disconnect.

17.2 The unit substation

It is often the case that a transmitter system of moderate average-primary-power demand (typically less than 500 kVA) must be integrated with a facility where the primary-power source is rated at a level many times greater, in the multiple-MVA category, for instance. If, in addition, the source voltage is considerably greater than the optimum for the transmitter system, a unit substation should be considered as part of the overall transmitter design.

A typical such substation is shown in Fig. 17-2. Such assemblies are intended to be used indoors. The key component of such a substation is the step-down line transformer. (Rarely does the voltage ever need to be stepped up.) It should be sized in VA rating as closely as possible to the maximum demand of the load system. The internal series impedance of this transformer, made up mostly of its leakage reactance, is the circuit element that will limit ultimate short-circuit cur-




Figure 17-2. A typical unit substation for loads up to 500 kVA.

rent. A typical value of internal impedance normalized to a load impedance into which full-rated VA would be delivered is 6% to 7%. This would limit ultimate symmetrical short-circuit current to approximately 16 times normal full-load current. The designer should consider that there may be other current-limiting impedances in series with the transformer. For instance, if the fault occurs close enough to the load end of the supply chain, the leakage reactance of the high-voltage rectifier transformer would impede conduction, which would further reduce the magnitude of the fault current. However, in any case the ultimate fault current will always occur in response to a short-circuit located immediately on the load side of the substation. Such a fault can never be considered impossible.

In situations where there is a mismatch between source and load VA ratings but there is no necessity for voltage transformation, the role of the step-down transformer can often be more economically performed by line inductors or reactors that are sized to have the same inductive reactance as the leakage inductance of the transformer. Often such reactors do not even require magnetic cores and, therefore, will never saturate. The windings must be sized and braced, however, for the steady-state and overload currents and the magnetic forces that they can produce.

The substation is also the natural location for the fused-manual disconnect, the circuit breakers, and the metering point for line voltage and current. In the case of an ungrounded delta output, it is also the ideal location for indicators for three-phase line-to-ground voltage balance. In addition to accommodating the main circuit breaker, the substation can usually house the branch circuit breakers that feed and protect the multiple subassemblies of a complete transmitter installation.

18. Switch-Mode Electronic Power Conditioning

These days, hardly any low-voltage dc power-supply designs are not “switch-mode” (or “switched-mode,” as grammar sticklers insist on calling them). From our previous discussions of pulse modulators we should be prepared to expect that the active devices in a switch-mode power supply are operated as electronic switches and not as linear or proportional amplifiers, and this is indeed the case. Furthermore, in modern usage the term switch-mode has become all but synonymous with solid-state, which is not always the case.

Why, we might ask, put a switchlike device in a dc power supply in the first place (other than to turn it on or off)? Consider the problem of converting the output of a source of primary power that can produce only dc, such as a battery or a photovoltaic solar cell, to a different voltage level. What we need, of course, is a dc transformer. Unfortunately, as we all know, such a thing does not exist. Or does it?

18.1 Switch-mode dc variable-voltage circuits (dc-to-dc converters)

Faced with the problem of converting the voltage of a dc source to a lower voltage at a load, a very literal technologist might reason that if a switch were placed between the source and the load, and that switch was repetitively turned on and off, the time-averaged voltage measured at the load would be smaller than that of the source simply because the source voltage was not being applied to the load all of the time. When the switch is closed, the load voltage is equal to the source voltage. When the switch is open, the load voltage is zero, as shown in Fig. 18-1a. Even though this could hardly be called dc transformation, the literalist is, of course, basically correct.

Suppose, however, that some additional components were added to the simple series-switch circuit—an inductor, a diode, and a capacitor—and they were arranged as shown in Fig. 18-1b. When the switch is closed, current will first start to build up linearly with time at an initial rate of V_{dc}/L and will exponentially approach a value of V_{dc}/R if the switch remains closed for a long-enough time interval. However, if the switch is repetitively opened and closed at a repetition rate that is high enough so that the interpulse interval, T , as shown in Fig. 18-1a, is small with respect to the circuit time-constant, L/R , then the time-averaged voltage at the load side of the switch will be $V_{dc} \times t/T$, where t is the length of time in each switching interval that the switch is closed. The current from the source will be discontinuous because it is being chopped by the switch. The current in the load, however, will be continuous, because the current in the inductor has no time to change from one switch-conduction interval to the next. When inductor current is not being supplied by the source through the closed switch, it is flowing through the shunt diode, which is often called a free-wheeling diode.

The circuit of Fig. 18-1b is called a single-quadrant chopper, or a buck regulator, because its output voltage is always smaller than the source voltage (not because they cost a dollar, which almost none do). When $t = T$, the switch is




Figure 18-1. Different types of switch-mode dc variable-voltage circuits.

closed all of the time and $V_0 = V_{dc}$. When $t = 0$, the switch is never closed and $V_0 = 0$. When t equals $1/2 T$, $V_0 = 1/2 V_{dc}$. If the output voltage is sampled and used as an input to a feedback circuit that can modulate the switch-conduction interval, t , the output voltage can be precisely regulated. This is called a pulse-

width-modulated (PWM) regulator, which usually runs at a precise clock rate ($1/T$). The higher this clock rate, the smaller the value of L can be and the smaller its core can be, if one is needed.

For switch-mode power conditioning, a full-control electronic switch is required, just as in the hard-tube pulse modulator. In general, the faster the switch can be turned on and off, the better. The efficiency of the circuit is primarily determined by the conduction-voltage drop of the switch, although there are "switching" losses as well that relate to circuit shunt capacitance. These losses are proportional to the clock rate. If the conduction-voltage drop of the switch approaches zero, the efficiency of the circuit approaches 100%. There are no intentional losses in the circuit. All of the power that leaves the source is delivered to the load. The peak current from the source is always equal to the average current in the load. If, for instance, the regulator is operating with $t = T/2$, the load voltage will be $V_{dc}/2$ and the load current $V_{dc}/2R$, which is also the peak source current. The load power is $V_{dc}^2/4R$. The peak source power is $V_{dc}^2/2R$, but for only half of the time. So the average source power is $V_{dc}^2/4R$, which is the same as the load average power, even though the source voltage is twice the load voltage.

The transient response of the circuit, which is related to its closed-loop bandwidth, also depends on the value of L/R , which determines how rapidly current can change in the load circuit. Note that none of the currents in the circuit ever change direction, or alternate. They are all dc. Is this, then, a dc transformer?

The buck circuit is not the only configuration of the key components that will produce the equivalent of dc transformation. The circuit of Fig. 18-1c is the dual of the buck circuit. It can only boost. When the switch is open all of the time, the output voltage is equal to the input voltage. When the switch is repetitively opened and closed, the output voltage is greater than the input voltage. Note that unlike the buck regulator, this circuit cannot operate with a unity switch-conduction duty factor. With the switch closed all of the time, the source will eventually become short-circuited by the dc resistance of the inductor and there will be no output at all.

However, at some switch duty factor less than unity, closing the switch momentarily will cause current to build up in the inductor at a rate of V_{dc}/L . When the switch is opened again, there will be energy stored in the magnetic field of the inductor of $W = 1/2 LI^2$, and the voltage at the load end of the inductor will rise to whatever value is required to cause conduction through the diode and the load circuit. (So far, this circuit is the same as the inductive-storage control-grid modulator shown back in Fig. 12-10.) If the switch remains open long enough for all of the inductor's stored energy to be delivered to the load circuit during each switch cycle, the average power in the load will be $P_0 = W \times f$, where f is the recurrence, or clock rate, of the switch closures. Assuming that V_0 is considerably greater than V_{dc} , the load power is $V_0^2/R = W \times f = 1/2 LI^2 \times f$, therefore

$$V_0 = \sqrt{P_0 \times R} = \sqrt{\frac{1}{2} LI^2 \times f \times R} .$$

Note that the output voltage is proportional to \sqrt{R} , other conditions being equal.




Figure 18-2. The single-quadrant dc-ac inverter with rectifier and pulse-width modulation regulation.

The inductor serves as a “charge pump,” lifting charge and energy from the source to a load whose voltage can be many times greater.

Figure 18-1d shows yet another arrangement of the key components. This one yields a buck/boost converter circuit that is capable of producing output voltage from zero to infinity—theoretically, anyway. The output voltage is zero at zero switch-conduction duty factor. As the switch-conduction duty factor increases, output voltage will increase in a fashion similar to that of the boost converter, except that the polarity of the load voltage is the opposite of that of the source. Unity switch-conduction duty factor cannot be reached because it will short-circuit the source through the dc resistance of the inductor. Note that even though the direction of current flow through the load is the opposite of that through the source, neither change direction. They are both dc, one intermittent and the other continuous.

18.2 DC-AC inverters

Although the circuits of Fig. 18-1, much like a transformer, are capable of converting one dc voltage level to another with theoretically no loss in power transferred, they are actually like the single-winding autotransformer in that one terminal is common to source and load. To realize nearly complete power transfer and the isolation of a true four-terminal network, there is no substitute for the

two-winding transformer.

The simplest dc-dc converter circuit with output-circuit isolation is the type shown in Fig. 18-2. This circuit uses a single full-control switch in a role similar to the switch tube in the transformer-coupled hard-tube modulator shown in Fig. 10-3. When the PWM switch is gated into conduction, current flows from the source, V_{dc} , through the primary winding of the transformer. If the switch is nearly perfect—that is, if it has low collector-emitter voltage drop during conduction—the voltage across the transformer winding will be nearly V_{dc} . During the conduction interval, current will also flow through the rectifier diode and load circuit connected to the secondary winding. The value of V_{dc} , the transformer primary/secondary turns ratio, and the resistance of the load will all determine the magnitude of the flat-top portion of primary current. Magnetizing current will also linearly rise throughout the conduction interval, producing an upwardly sloping total primary current.

The currents in the primary and secondary windings are unidirectional, or dc, but they are not continuous. However, there can be no average value of voltage across any transformer winding because the windings are, in the steady-state, simply lengths of wire. The size of the core they are wound around makes no difference. In order to satisfy the criterion that there is no average voltage, the volt-time products must be the same in both polarities. The energy stored in the transformer magnetizing inductance will cause the transformer voltage to reverse the instant the PWM power switch is turned off. The current that had built up in the magnetizing inductance during the conduction interval wants to continue flowing, just as it did in the transformer-coupled hard-tube modulator with a diode load. Rather than being dissipated, the energy stored can be returned to the power source by means of an additional transformer winding, a clamp winding, and a clamp diode. If the clamp winding has the same number of turns as the primary but is of opposite polarity as shown, voltage will reverse across the clamp winding at the same time it reverses across the primary. The clamp diode will conduct as soon as the clamp winding voltage becomes more positive than the input voltage, V_{dc} , and the magnetizing current will continue to flow back into the source, linearly decreasing with time, until the energy has been restored to the source. The voltage across the switch during the non-conducting interval will be clamped at twice V_{dc} . The switch duty factor cannot exceed 50%, because that would produce the limiting case of equal volt-time products. If the clamp winding had a different number of turns, higher peak voltage on the switch could be exchanged for a higher duty factor, or vice versa.

The load voltage will be a function of the switch duty factor. At constant clock frequency, or constant value of interpulse interval, T , the duty factor will be proportional to the on time, t , or pulse width. Modulation of pulse width in response to an error signal generated by the difference between the sampled output voltage and a reference voltage can result in a regulated output voltage. The greater the clock rate, the smaller the transformer volt-time products will be and the smaller the transformer core can be made before saturation occurs. When attempting to produce high-voltage output, however, there are limits to the optimum switching rate.

The circuit just described is called a single-quadrant forward converter be-

cause current flow in the load circuit and in the switch occurs at the same time. A circuit similar to it, called a flyback converter, operates in the opposite mode. The polarity of the secondary is reversed. During switch conduction, energy is deliberately stored in the transformer magnetizing inductance because only magnetizing current flows. When the switch is turned off, the energy stored in the magnetizing inductance is transferred to the load (in similar fashion to the boost and buck/boost converters, which also operate in the flyback mode). No clamp winding is required, of course. The most common example of the flyback converter is the high-voltage power supply of a television receiver, which makes use of the energy stored in the horizontal-deflection coils during the retrace intervals at the ends of each horizontal trace.

In transmitter circuits, even low-power ones, the capabilities of the single-switch converter are usually inadequate. Moreover, the current flow in both primary and secondary windings is unidirectional, so full use is not made of core magnetization, meaning that larger cores are needed for a given power level for this kind of circuit than for others. A commonly used alternative is the push-pull converter, shown in Fig. 18-3. It uses two switches that are connected to opposite ends of a center-tapped primary winding. The switches conduct alternately for equal times. The currents in the two halves of the primary winding are unidirectional, but they magnetize the core alternately in opposite directions, so there is no dc component of magnetization. The rectifier circuit shown is a full-wave center-tapped type. It also conducts alternately in synchrony with switch conduction. (Note that any type of full-wave, single-phase rectifier would be appropriate.) No free-wheeling diode is required for the filter inductor, because total rectifier conduction is continuous, so there is always a path for filter-inductor current.

In the circuit shown, the transformer-core magnetization capability is fully utilized, but the ohmic capacity of the windings is not. Each winding conducts half of the average load current. But it is in the form of a square wave whose peak current is equal to the average load current of the 50% duty factor (normalized to the transformer turns ratio). Remember from Fig. 6-2 that the average value of this waveform is $I_{PK}/2$, but the RMS value is $I_{PK}/\sqrt{2}$. Therefore, the




Figure 18-3. The push-pull switch-mode dc-ac inverter with full-wave, center-tapped rectifier.




Figure 18-4. The half-bridge, or "totem-pole," switch-mode dc-ac inverter.

dissipation in the winding is twice the minimum achievable value. The same is true of the secondary winding for the type of rectifier shown. Despite this drawback, however, such a circuit may prove useful in certain circumstances. The push-pull inverter produces a transformer input of ac—or, more correctly, av (alternating voltage)—that has a peak-to-peak value of $2 \times V_{dc}$.

An alternative type of inverter, shown in Fig. 18-4, produces true ac input to the rectifier-transformer primary. It is called the half-bridge, or "totem-pole," inverter, and it is a circuit topology that we have certainly seen before. Like the push-pull inverter, it uses two switches that are alternately gated into conduction. In order to achieve an input voltage to the transformer of $2 \times V_{dc}$, the dc source must be $2V_{dc}$ or, as shown, plus and minus $1V_{dc}$. In the push-pull inverter, the only anomaly capable of producing a dc component of transformer magnetization (assuming a perfectly center-tapped primary winding) is to time the push-pull switch conduction at some time interval other than 50/50. In the half-bridge inverter, a coupling capacitor in series with the primary winding will assure that there is no dc component. Without the capacitor, differences between the absolute values of $+V_{dc}$ and $-V_{dc}$, and a switch timing other than 50/50 can both produce a dc offset. If a source voltage of $2V_{dc}$ is used, there is no alternative to the coupling capacitor. It will automatically assume an average voltage drop nominally equal to V_{dc} , but it will differ from that amount by whatever increment is required to assure that the inverter output is true ac with precisely equal positive and negative volt-time products.

Although the half-bridge circuit has been shown as a generic inverter, it would not necessarily be a component of a dc-dc converter that might be used to drive an arbitrary load. Unlike a rectifier load, which has a unity power factor with square-wave input, the current and voltage of an arbitrary load cannot be expected to have the same wave shape, let alone change direction at the same

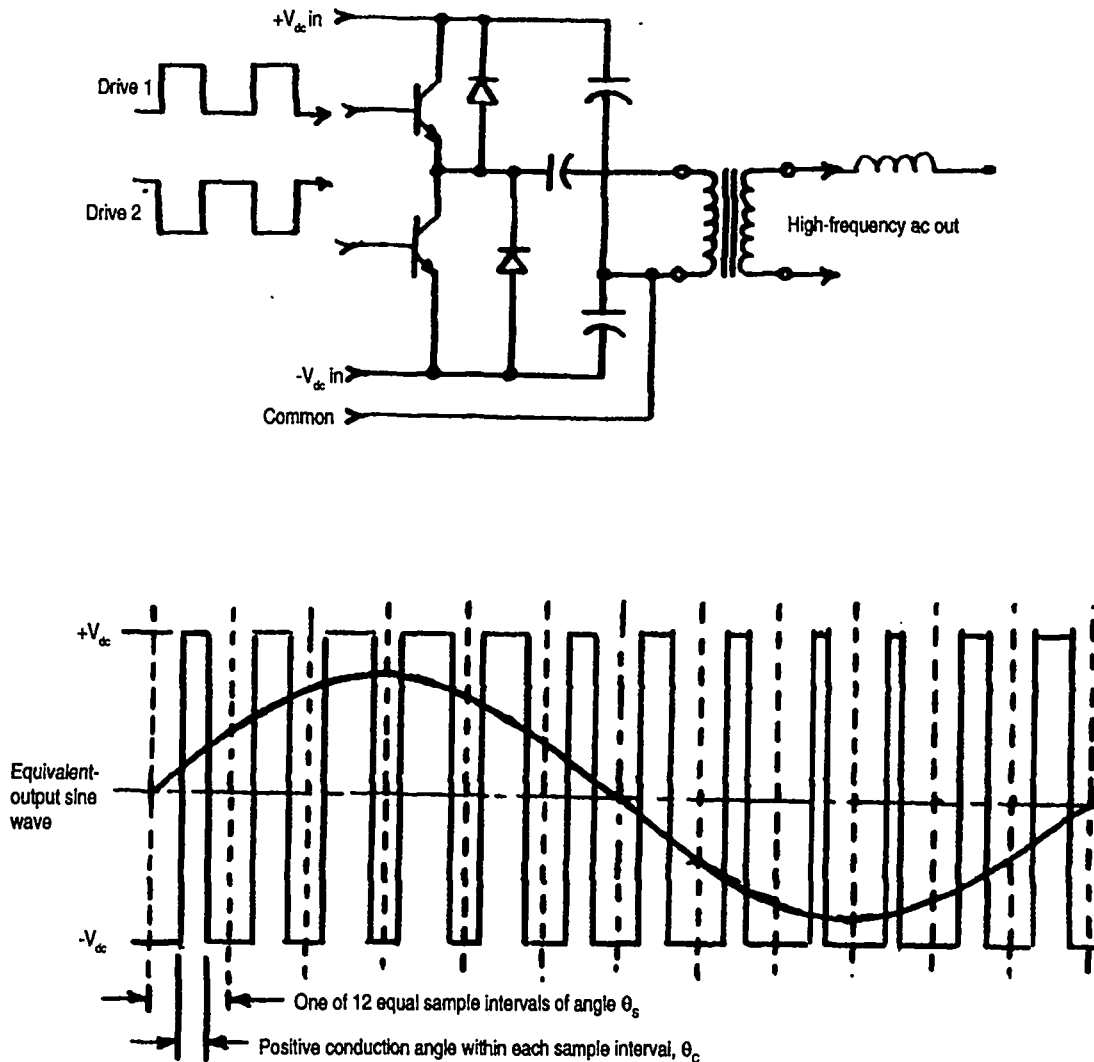


Figure 18-5. Pulse-width-modulation strategy for the bridge inverter to synthesize sine-wave output.

instant. For that reason, the generalized inverter must have the internal capability to pass currents in either direction at all times. This is the function of the anti-parallel diodes that shunt the inverter switches.

Simultaneous conduction, or shoot-through, of the switches in the half-bridge is always a possibility and can cause serious stress on the switching devices. A gating strategy is required to assure that one switch has ceased conduction before the other one begins. The power MOSFET requires less attention in this regard than the power bipolar transistor, which has inherent turn-off delay due to minority-carrier cleanup. (In the push-pull inverter, the magnetizing inductance of the transformer will at least slow the rate-of-rise of shoot-through current, making the problem less severe.)

Figure 18-5 shows a half-bridge inverter being used to synthesize a sinusoidal output waveform by means of a pulse-width-modulated bipolar output. The inverter operates at a clock frequency that is many times greater than the output frequency. In the example shown, it is 12 times as great, dividing up the output sinusoid into 12 sample intervals, each one dealing with the appropriate 30°

segment of the output wave. The load has an inductor-input band-pass filter. The rail-to-rail bipolar input voltage is developed across this inductor. The proper half-angle of conduction is different for each sample interval. It is determined by solving the following equation

$$\frac{\sin \theta_c}{2} = \frac{1}{4 \sin \left(\theta_1 + \frac{\theta_s}{2} \right)} \left\{ \cos \theta_1 - \cos(\theta_1 + \theta_s) + \frac{K}{2} \left[\theta_s + \sin \theta_1 \cos \theta_1 - \sin(\theta_1 + \theta_s) \cos(\theta_1 + \theta_s) \right] \right\}$$

where θ_s is the angle of the equal-sample interval, θ_c is the positive conduction angle within each sample interval, θ_1 is the angle (referred to the output sine wave) at the beginning of each sample interval, and K is the ratio of peak output sine wave to input V_{dc} .

Just such an inverter was designed and built, but it used triode vacuum tubes instead of transistors as switches. The triodes were driven from a computer program that approximated the solution of the equation to produce frequency-shift-keyed, extremely low-frequency (ELF) signals. This application was the prototype for an early version of the ELF submarine communication system.

With more switches, even more can be done with an inverter. The full-bridge inverter shown in Fig. 18-6 uses four switches, two of which conduct simultaneously (1A and 1B, or 2A and 2B). In return for the additional complexity and switch dissipation, the full-bridge inverter produces a peak-to-peak input to the transformer primary of $2 \times V_{dc}$ for a source voltage of V_{dc} . A coupling capacitor in series with the transformer primary once again can insure that there is no dc component by automatically assuming whatever average voltage across it is re-




Figure 18-6. The full-bridge switch-mode dc-ac inverter.

quired to nullify any timing asymmetry. Primary voltage asymmetry is unlikely because the dc source is first connected across the primary in one direction (1A and 1B conducting) and then in the other (2A and 2B conducting). A three-state, pulse-width-modulated output voltage is achieved by varying the relative timing of square-wave gating signals, as shown. During the zero, or neutral, state, both transformer primary terminals are clamped either to the source high-side or to the return (either 1A and 2A conducting or 1B and 2B). This clamping short-circuits the primary winding during the neutral state, minimizing possible ringing or other transient effects when zero output is desired.

The switch-mode converter and inverter circuit topologies discussed are by no means the only ones. Resonant- and quasi-resonant-type inverters have been mentioned before (see Fig. 9-45) and will be again, but even their inclusion does not round out the complete list. The most modern switch-mode circuits have been made practical by the development of powerful solid-state, full-control switches such as the power MOSFET. (In order for the switch function to be performed by a single device, it must have full-control properties. In the past, large-power circuits were switched with half-control devices such as SCRs and even Ignitrons, that had to be shut off by "forced" commutation, which required additional components and imposed a serious limit to the switching rate.)

As an interesting side note, let me say that had MOSFETs been in existence in the early days of electricity, it is conceivable that the point of view of Thomas Edison, who was an influential proponent of dc power distribution, might have prevailed over that of George Westinghouse, the most powerful proponent of ac power distribution. (At least the controversy would have been more interesting, especially considering that the highest-voltage power transmission system in this country is the Pacific Intertie, which operates at 1 MVdc and uses power rectifiers and power inverters at its terminals.)

Even before the advent of modern super-power solid-state switches and rectifiers, dc-ac inversion and dc-dc conversion was common. Literally millions of automobile radio receivers used vacuum-tube amplifiers with anode supply voltages up to 300 Vdc. They were powered from automotive storage batteries whose voltage was as low as 6.3 V at one time. The repetitive switch function in these radios was performed by a device called a vibrator, in which contacts of an actual mechanical switch were magnetically opened and closed at a low audio frequency (low enough that one could tell whether or not it was vibrating by the sound it made). The most sophisticated of the vibrators was a "synchronous vibrator." It had two sets of contacts, one in the primary circuit of a step-up transformer and the other in the secondary to directly rectify the output. Switch-mode power conversion, therefore, predates the solid-state revolution by a number of decades.

Up until now, the discussion of converters and inverters has been limited to those that operate from a dc power source. Conversion to a different voltage level with no power loss or inversion to ac requires a controllable switch or switches. There is no alternative. The faster the switch or switches can be made to turn on and off, the smaller the other components can usually be made, especially the iron-core components. For transmitter applications, however, these circuits may be required to produce relatively high voltages, or at least voltages

that are many times greater than that of a typical dc source. As output-voltage levels increase, step-up transformers require more insulation between primary and secondary windings. More insulation means greater physical separation between windings and greater leakage inductance. The reactance of this leakage inductance increases with increasing frequency, thus limiting the amount of voltage step-up that can be achieved for a given output current. For each specific application, there will exist an optimum transformer design and inverter operating frequency that will be defined by the source voltage and the requirements for output voltage and power. Although optimization studies are rarely performed, it is true that few high-power and high-voltage systems using switch-mode, high-frequency technology, operate at more than 10 kHz. And the true "optimum" frequency in many cases may be considerably less. As one representative of a transformer and power-supply vendor succinctly put it, "High frequency and high voltage don't mix."

When the source of primary power is ac to begin with, or, as is sometimes the case, when the prime mover is something as non-electrical as a rotating shaft, the choice of switch-mode power conditioning should deserve even more thought. The switch-mode power supply operating from an ac source is often called an off-line switcher, meaning that it operates "off" a power line, not that it is broken and therefore "off-line." The first thing that must be done in an "off-line" power conditioner is rectification to produce the internal dc rail for the high-frequency inverter, whose output is then transformed in voltage and eventually rectified again to get back to dc. At low voltages and relatively low power levels, all of these steps may be justifiable for a high-frequency inverter in terms of overall size, weight, and efficiency, especially if output-voltage variability and regulation are required. For this kind of device, "high" efficiency is usually in the 65-85% range, which is not high at all when compared with the overall efficiency of a high-voltage, high-power transformer-rectifier system. These systems can easily exceed 95% efficiency. But the size and weight of such a high-voltage power supply operating from commercial 60-Hz ac source can be immense when compared to one that operates from a much higher-frequency ac source.

Techniques have been developed that offer dramatic reduction in the size and weight of iron-core components operating even from 60-Hz ac. They involve the use of wire that may be one-tenth the diameter of that used to wind a transformer of conventional size and efficiency. Obviously, such a ploy increases winding losses, but the windings must then be cooled effectively to keep them from overheating. The transformer efficiency will also be degraded by using smaller wires, perhaps to 90%. In return, however, the size and weight of the core can be drastically reduced because the same number of winding turns can fit in a much smaller core window. The overall efficiency may be not be much greater than that of a complete off-line high-frequency, switch-mode power supply, but the circuit complexity might be two orders of magnitude less. And it might be less expensive and more reliable as well. Such techniques are only rarely investigated now because of the "technological correctness" of the electronic alternative, but high-frequency, switch-mode power conditioners are neither simple nor cheap. If a designer has the time, all options should be studied.

When the source of system input power is the prime-mover itself, and the

choice of mechanical-to-electrical power conversion is at the system architect's discretion, even more options are available. High-frequency three-phase alternators operating up to 1 kHz are practical. It goes without saying that 400-Hz alternators are practical because they are standard for most large-power aircraft auxiliary power units. Iron-core components can be made quite small and lightweight to operate from 400 Hz—even without resorting to the ultra-miniaturization techniques mentioned above—and they can be made even smaller to operate at 1000 Hz. The problems of leakage reactance and winding skin-effect are only moderate. If an alternator is dedicated to a single power supply, voltage regulation and variability can be accomplished by modulation of the alternator-field excitation. The alternator can also perform the function of high-speed electrical disconnect.

18.3 Voltage-multiplier rectifier circuits

If, even after visiting all of the aforementioned considerations, the designer finds the lure of the high-frequency, switch-mode conditioner still irresistible, then he or she should know that there are ways to overcome the limited transformer step-up ratio. One of them is the voltage-multiplier rectifier circuit. A basic form of a voltage multiplier rectifier is shown in Fig. 18-7. It is an iterative circuit that can be expanded to almost any level of voltage multiplication. Illustrated are the two-pulse voltage doubler, the three-pulse voltage tripler, and the four-pulse voltage quadrupler.

The "pulses" referred to are, in the cases illustrated, polarity reversals of the input alternating voltage. During the first "pulse," the high side of the transformer secondary will be positive with respect to the grounded return. Current will flow clockwise through the first capacitor and the first diode. If the source impedance is low enough, the capacitor will be charged to the peak value of the input voltage, V , especially if the wave shape is rectangular. When the polarity of the input reverses, assuming that the second capacitor has no charge as yet, the total voltage in the loop, which includes the first two capacitors and the second diode, will be $2V$. Current will flow counterclockwise. If the second capacitor is many times smaller than the first, it will be charged to a voltage of $2V$ with the polarity as shown. During the third "pulse," the source voltage reverses polarity again. The first and the third diodes will now conduct. Current I_1 flows through the first diode and replenishes the charge on the first capacitor. If the first capacitor has lost no charge, I_1 will be zero because there will be no net loop voltage. The net voltage in the loop, including the first three capacitors and diodes, will be $2V$ again and be supplied by the voltage across the second capacitor. The source voltage and the voltage across the first capacitor are of opposite polarity and cancel each other. Assuming that the third capacitor is much smaller than the second, current I_2 will charge the third capacitor to $2V$ with the polarity shown. The fourth "pulse" will produce current through the first and fourth diodes and all of the capacitors. If, once again, the fourth capacitor is much smaller than the third, it will be charged to $2V$, which is the initial loop voltage ($2V + V + V - 2V = 2V$). As described, the doubler will have an output voltage of $2V$ after two "pulses," the tripler will have an output of $3V$ after three "pulses," and the quadrupler will have an output of $4V$ after four "pulses."

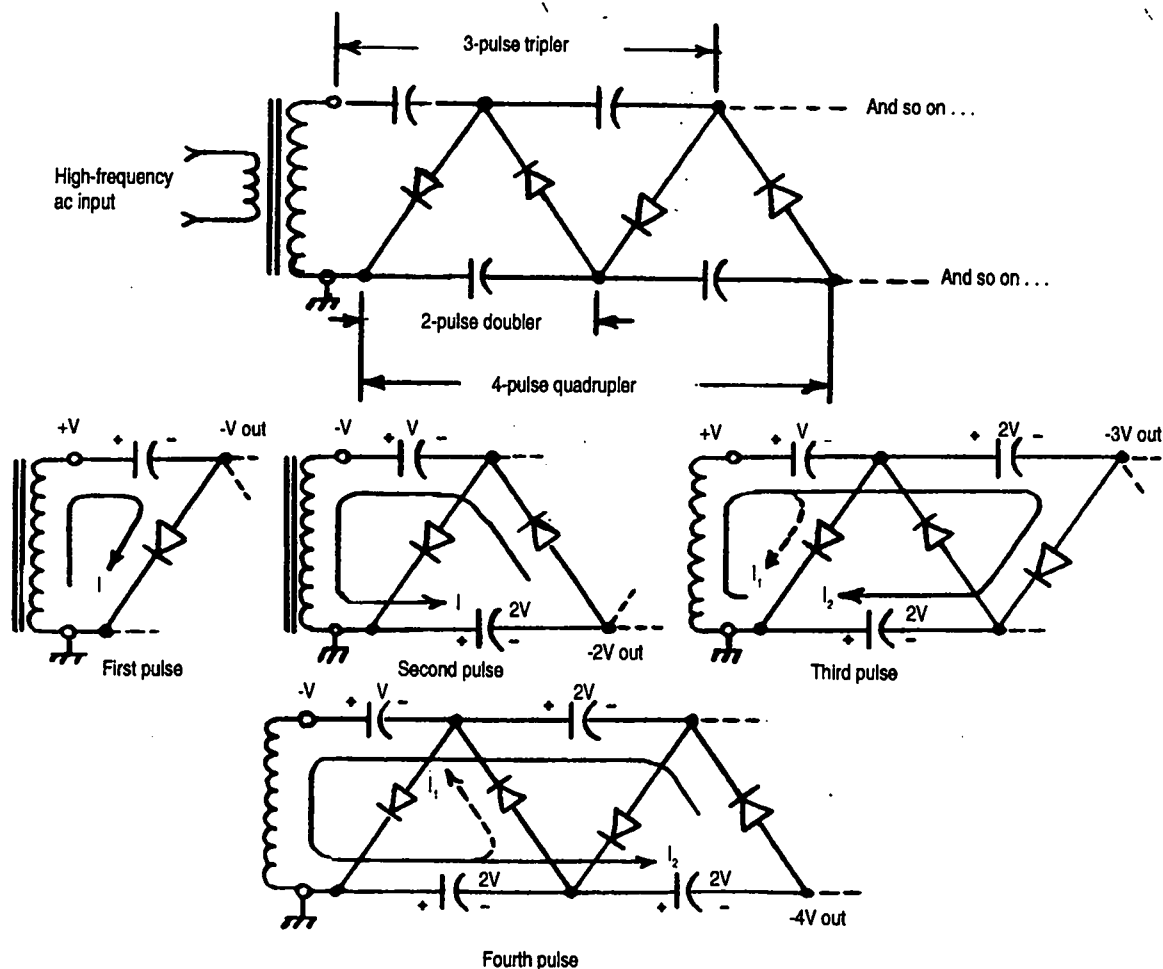


Figure 18-7. Multiple-phase voltage-multiplier type of rectifier circuits.

The theoretical multipliers described are less than practical because the successive capacitors must be much smaller than the preceding ones. Practical multipliers use nearly identical capacitors in each stage. Many more input "pulses" are required to fully charge all of the capacitors because of the voltage lost by early capacitors in supplying charge to later capacitors. In addition, there is a significant time lag between the time when the input voltage changes and when the output changes. This lag affects the bandwidth and transient response of a closed-loop regulation system. There are virtually as many types of voltage-multiplier rectifier circuits as there are conventional rectifiers, including full-wave center-tap, full-wave bridge, etc.

Voltage multipliers need not be restricted to single-phase input. Figure 18-8 shows a three-phase voltage-multiplier rectifier that is driven by three square-wave half-bridge inverters whose timing is such that the outputs of each are delayed by $1/3$ of a master-clock period, thus producing the effect of three-phase ac. (There is no reason to be limited to three phases either, or multiples thereof; four- or five-phase systems are no less practical.) As previously noted, one of the advantages of polyphase rectification of ac having a sinusoidal waveform is that conduction is restricted to the peaks of the sinusoids, and the smaller the conduc-




Figure 18-8. Three-phase voltage-multiplier rectifier circuit.

tion angle, the closer the unfiltered rectifier output is to pure non-varying dc. When the ac is produced by square-wave inverters, however, there is no such advantage to a polyphase system, because the tops of the ac input are nominally flat. The polyphase advantage, therefore, results from dividing up the total average current into multiple parallel, but time-sequential paths, that time-share the load.

In the practical example shown, the outputs of the inverters are connected to the primary windings of a wye-connected transformer, the "neutral" of which is connected to the common point, or ground, of the bipolar dc source, $+V_{dc}$ and $-V_{dc}$. Given the timing relationships shown, and starting with phase A positive (switch A1 conducting), we see that phase B is negative (switch B2 conducting). Secondary winding 1 will be positive by an amount proportional to V_{dc} , and secondary winding 2 will be negative, also by an amount proportional to V_{dc} . The actual voltages will depend upon the transformer turns ratio and other transformer and load-current factors. Conduction through diode D2 will charge capacitor C2 to an amount proportional to $2V_{dc}$, in the polarity shown. The con-

duction interval will have a duration of $1/3$ of a master-clock period. For the next $1/3$ period, phase *B* will be positive and phase *C* negative (switches *B1* and *C2* conducting). Capacitor *C3* will charge to a voltage proportional to $2V_{dc}$ through *D3*. For the final $1/3$ period, phase *C* is positive and phase *A* negative (switches *C1* and *A2* conducting), and *C1* will charge through *D1*. As the cycle continues, charge accumulated by *C1* will be shared with *C5* through *D5*; charge on *C2* will be shared with *C6* through *D6*; and charge on *C3* will be shared with *C4* through *D4* until all capacitors are charged to voltages proportional to $2V_{dc}$. The output capacitor *C7* will eventually charge to a voltage proportional to $V_{dc}(2n + 1)$, where n is the number of multiplier stages, or series capacitors, through summing diodes *D7*, *D8*, and *D9*. The summing diodes charge the output capacitor to a voltage that is greater than the sum of the series-capacitor voltages by an amount proportional to V_{dc} . In the case illustrated, $n = 2$, so the output voltage will be proportional to $5V_{dc}$.

The input-output characteristics of the voltage-multiplier rectifier are not unlike those of a transformer-rectifier system having a high transformer step-up ratio. The dc output voltage of the multiplier-rectifier can be many times the peak voltage of the ac source. In return, we can expect that the source current will be higher than the load current by the same factor—or more, if there are circuit power losses—simply by law of the conservation of energy. The same result is reached by recognizing that the capacitors are in series with respect to the load but in parallel with respect to the source. Charge delivered to the load over any period of time is removed from each capacitor simultaneously, but the charge must be replaced individually from the source. If charge Q is removed and there are n capacitor stages, a total charge of nQ must be supplied by the source over the same period of time to maintain steady-state equilibrium.

One particular form of the voltage-multiplier rectifier is often used to produce very high voltages in particle-accelerator applications. It is called the Cockroft-Walton generator, a type of which is shown in Fig. 18-9. This rectifier uses a full-wave center-tapped voltage multiplier that is formed by superimposing a mirror image on the half-wave type shown in Fig. 18-7 to produce a voltage step-up of 10. For this type of voltage-multiplier rectifier, output voltages up to one megavolt are common. At these high voltage levels, such generators have an inherent physical superiority over simple, high-voltage transformer-rectifier assemblies: insulation. Both types use about the same number of individual rectifier elements, but in the Cockroft-Walton generator the stages can be spaced and graded (using deck corona shields of adequate radius of curvature) so that air can provide sufficient dielectric strength even at such high voltages. The type illustrated can be of prodigious size. Even so, with a voltage step-up of 10, the peak input voltage must be 100 kV for an output voltage of 1 MV. Such generators typically operate at source frequencies of up to 10 kHz, with the input supplied by high-power audio amplifiers that use power vacuum tubes as output devices. For an average-output current of 10 mA at 1 MV, the average power input must be at least 10 kW (and the reactive VA may even be higher).

18.4 High-voltage dc power supply using a quasi-resonant inverter

The inverters discussed so far have had square-wave outputs, or rectangular-pulse outputs if it is pulse-width modulated. These waveforms require the abrupt "chopping" of a dc input. This activity introduces significant repetitive transient pollution that must be suppressed, and it produces corresponding stress on the electronic switches that must do the chopping. The switches, moreover, must either have full-control capability or they must be turned off by "forced-commutation" involving additional components and switching-speed limitations. The quasi-resonant inverter, which was introduced in Fig. 9-45, does not share these problems. A high-voltage, off-line, dc power supply that uses such an inverter along with a high-step-up multiplier-rectifier, is shown in Fig. 18-10.




Figure 18-9. The Cockcroft-Walton generator for very high voltages.




Figure 18-10. High-voltage switch-mode power supply using resonant-type inverter and full-wave bridge-type voltage-multiplier rectifier.

A dc source is created by direct three-phase, full-wave (six-pulse) rectification of the line, producing two balanced dc rails. A neutral return for the half-bridge inverter is provided by a center-tapped capacitor filter at the rectifier output. Power transfer from the source to the load occurs during the half-cycle conduction intervals of the two SCR power switches, *SCR1* and *SCR2*. The conduction path for *SCR1* is in the positive direction through *L1*, *C*, and the rectifier-transformer primary winding. The conduction path for *SCR2* is in the negative direction through *L2*, *C*, and the transformer primary winding. Inductors *L1* and *L2* are of equal inductance. The values of *L1*, *L2*, and *C* are chosen to provide two important properties of the overall circuit. First, the characteristic impedance of the discharge circuit ($\sqrt{L/C}$), where *L* is either *L1* or *L2*, is chosen so that it is considerably smaller than the value of resistance that is reflected back from the rectifier and load into the transformer primary. Therefore, the series-resonant circuit is underdamped and will ring. The current will have a damped sine-wave shape. When the current through *SCR1* passes through zero, *SCR1* will be automatically commutated off. The negative-going half-cycle following *SCR1* conduction will flow through *D1*. When current reaches zero again, it will stop until *SCR1* is gated on again. The same sequence occurs in the opposite direction with *SCR2* and *D2*. The values of *L* and *C* also determine the period, or oscillation frequency, of the discharge current. The parameter that allows variability (or regulation) of output voltage is the controllable time interval, Δt , between each power-transfer cycle. Even though the transformer primary currents are sinusoi-

dal segments, the primary voltage is "rectangularish" because of the hold-up produced by the reflected load-circuit capacitance.

Two significant differences are apparent between the square-wave and the quasi-resonant inverter. The first is that the peak switch current for the quasi-resonant inverter is several times greater than the load current (typically 5 to 6 times as great). This is necessary for the primary circuit to perform in underdamped, self-commutating fashion. This circulating current is not provided by the primary power source. It contributes to inefficiency only in that the ohmic losses in the inductor and transformer primary will be proportional to the square of the circulating current. On the other hand, the switch only needs to have half-control capability. It plays no role in determining either the amplitude or duration of the current, which, by virtue of its sinusoidal wave shape, has no sharp discontinuities, either. The second difference is that the quasi-resonant inverter, in order to have control over load voltage, operates at variable frequency rather than constant frequency and variable pulse duration.

The rectifier shown is the same full-wave, center-tapped type used in the Cockcroft-Walton generator. At least one practical version of this circuit uses a multiplier with a step-up ratio of 20, giving an output voltage of 75 kVdc with a transformer turns ratio of 21 and 480-V, three-phase input. Like the Cockcroft-Walton generator, the power supply is entirely air-insulated, thus taking advan-




Figure 18-11. Block diagram of traveling-wave-tube amplifier (TWTA) using high-frequency switch-mode electronic power conditioning (EPC).

tage of the large dimensional spacings possible with voltage-multiplier rectifiers.

18.5. The electronic high-frequency power conditioner as applied to a microwave tube

The most common microwave-tube application of the high-frequency electronic power conditioner (EPC), is in the traveling-wave-tube amplifier (TWTA), a generic block diagram of which is shown in Fig. 18-11. Such TWTAs abound in aircraft and satellite environments where size and weight are of primary concern. Many such EPCs operate from so-called 28-Vdc aircraft primary power (which has a specified minimum value of 19 Vdc). There is a practical upper limit to the amount of power that can be converted from such an input bus, somewhere around 1 kW. (Higher-power can be extracted if the input bus is 208-V, three-phase, 400-Hz because direct rectification will yield internal dc rail voltage of almost 300 Vdc, from which MOSFET inverter switches will operate more optimally, inverting higher power at higher efficiency.)

A TWTA will have no fewer than two independent, high-frequency inverters that operate at the highest practical clock rate, usually between 10 kHz and 100 kHz. The higher the voltage and power levels, the lower the clock rate is likely to be for overall optimum performance. The first inverter supplies the input to the filament and grid-bias transformer-rectifier circuit (and to the grid modulator, if it is the direct-coupled type). After this inverter has been turned on, there will often be a read-back of filament voltage and current and grid bias voltage. These values are usually summed as a single optically coupled status signal that is sent to the ground-level control electronics, which will include a filament time-delay circuit. The filament and grid-bias circuits float at the TWT cathode voltage, which is the output of the helix/cathode transformer-rectifier. This voltage can be anywhere from a few kilovolts to a few tens of kilovolts, but it is most usually in the 3-kV to 5-kV region. Its input comes from the second high-frequency inverter, which is turned on only after the requisite cathode-heating time-delay has elapsed. The helix/cathode or cathode/ground voltage is the highest and must be the most precisely controlled voltage in the system. A sample of this voltage is ordinarily used as the input to the PWM regulator for this inverter. The TWT used in such amplifiers will almost always have a depressed collector or several depressed collectors (as many as five stages of depressed collector have been featured in some high-power space-based TWTs). Each collector will have a separate transformer-rectifier, or at least a rectifier operated from a separate secondary winding that shares the same primary and core. In a two-stage depressed-collector TWT, as illustrated, the Collector-1 voltage may be from 80% to 100% of cathode-ground voltage (0 to 20% depression) and the Collector-2 voltage may be from 60% to 80% of cathode-ground voltage (20% to 40% depression). The cathode-ground rectifier may handle as little as 3% to 5% of the total beam current, with the remainder being split between the grid and the multiple collectors.

Some TWTs, especially those with coupled-cavity circuits, may be endangered if beam current flows before beam voltage reaches the nominal operating value during turn-on. (If the tube is pulse modulated, this is not an additional problem. Pulses to the modulator can be simply inhibited until beam voltage is at the




Figure 18-12. Practical TWTA switch-mode power conditioning using a buck regulator and push-pull inverter.

proper value.) A tube intended for continuous duty (CW) may still require some means of inhibiting beam current during turn-on. Often a modulating or isolated anode is provided, as shown. It is held at cathode potential by a series resistor. When cathode voltage reaches the proper value, the anode is connected to ground by means of a high-voltage but low-current relay, which turns on beam current.

A practical implementation of the high-voltage dc portion of an EPC in a TWTA is shown in Fig. 18-12. The power inverter is the push-pull type. Each MOSFET switch operates at 50%-duty factor. Output voltage control is accomplished by means of a buck-type switch-mode regulator that is directly in series with the dc input source. This source is pulse-width modulated to regulate or limit either the input voltage and current to the step-up transformer or the dc output voltage and current, depending upon which sets of sample points are connected to the dual error amplifiers. The error amplifiers compare the sampled signals with a reference voltage, thus developing an amplified error signal at the parallel-connected outputs. The error signal is compared with a sawtooth-shaped voltage in the pulse-width modulator called the comparison ramp. The ramp voltage always starts upwards at the same voltage and resets to that voltage at the end of each clock period. The buck regulator is turned on at the beginning of each clock period, coinciding with the start of the ramp. When the ramp voltage exceeds the amplified error voltage, the drive to the regulator is terminated. For instance, if the error signal is $E1$, the duration of the regulator drive will be $PW1$, if $E2$, it will be $PW2$, and if $E3$, it will be $PW3$, as illustrated. This process is often

called delta modulation, which converts voltage change to pulse-duration change.

The particular circuit described was used in a multi-purpose radar/communications transmitter. It utilized a high-performance TWT that was required to operate with low-Doppler sidelobes, which meant very-low-ripple sidebands. One clever trick used to achieve this requirement was to synchronize the inverter clock with a multiple of the radar repetition rate. This effectively "hid" the power-supply-inverter ripple components under the PRF lines of the pulse spectrum. At constant PRF, this technique would work fine. In this particular application, however, PRF was not constant. In fact, it changed frequently under computer control. In a conventional delta modulator, the ramp rate is constant. Therefore, if the clock rate is changed, there will be an immediate increase or decrease in buck-regulator duty factor, depending upon whether the system PRF went up or down. This change in the clock rate causes a corresponding transient increase or decrease in high-voltage dc output. The feedback will eventually correct the output error but not very quickly. The closed-loop bandwidth of such a regulator is only tens of hertz. In actual operation, the regulator was continually hunting for the proper output voltage because PRF changes happened with a frequency not much lower than the closed-loop bandwidth. The solution, which was successfully implemented, was to change the ramp rate as well as the clock rate, thus making the regulator duty factor independent (almost) of the clock rate.

The value of the output inductor, L , of the buck regulator must also be carefully chosen. Its purpose is to smooth the discontinuous output of the regulator. However, if it is too large and does the job too well, it severely limits the response time of the regulator loop—which isn't that fast to begin with. The current in the inductor will rise ($I1$) when the regulator is turned on and fall ($I2$) when the regulator is turned off. A peak-to-valley variation of 10% for this type of circuit is considered optimum by many designers.

18.6 An idea too clever not to share (It's not mine, by the way.)

As mentioned before, the leakage inductance of either a step-up or a step-down isolation transformer that is used in a switch-mode power supply imposes a fundamental limitation on the circuit's optimum switching frequency. However, there is a power-supply design covered by US patents that, instead of fighting leakage inductance, actually makes use of it in reaching switching frequencies of up to a megahertz. This design does not appear in high-voltage, or especially high-power, products, but the products it can be found in are nevertheless quite useful in many transmitter applications.

A simplified schematic diagram of the power-transfer portion of a typical circuit, along with its voltage and current waveforms, is shown in Fig. 18-13. The circuit comprises a single-ended MOSFET power switch; a transformer, T ; a charging diode, $D1$; a storage capacitor, C ; a free-wheeling diode, $D2$; a smoothing inductor, L ; and the useful load, shown as a resistor. The leakage inductance of the transformer, $L_{LEAKAGE}$, performs what should be a familiar role. Starting at the beginning of an energy-transfer cycle when the FET switch is first turned on, current will flow in the primary and secondary windings of the transformer and through diode $D1$. The current waveform through $D1$ will be in the form of a




Figure 18-13. Megahertz-rated quasi-resonant (dc-resonant charging) switch-mode power supply.

half sine-wave, the frequency and amplitude of which are determined by $L_{LEAKAGE}$ and *C*. Assuming that steady-state load-current conditions had already been reached, a steady value of current will flow in *L* and the load. This current is either replaced or maintained by the paths through either *D1*, *C*, or *D2*. As soon as the rising current through *D1* exceeds the constant value of load current, there will be an excess of current available. This current charges capacitor *C* in a dc-resonant fashion that is almost identical to the dc-resonant charging of the line-type modulator PFNs previously discussed. Except in this case, $L_{LEAKAGE}$ performs the role of the charging inductor in the line-type PFNs. But it should be noted that a conventional transformer may not have sufficient leakage inductance to optimize this circuit. Therefore, the transformer, as shown in the simplified cross-section drawing, must have primary and secondary windings that are deliberately separated from one another on a pot core by a dielectric spacer. This is done to decrease the coefficient of coupling, *k*, and increase the leakage inductance, which is proportional to $(1 - k^2)$. This strategy also enhances the circuit's voltage-hold-off capability and decreases the capacitance between input and output terminals, thus making the circuit more like a true four-terminal network and more useful in many transmitter-circuit applications requiring "floating" power sources.

Once the current through *D1* approaches the end of its half-sine-wave exist-

ence, it falls below the value of the load current. This terminates the resonant charging of C and starts its constant-current discharge. Before that begins, however, C has been charged to almost twice the value of the transformed source voltage through the dc-resonant charging action. The constant-discharge current is the load current, which is kept constant by L . The voltage across C will linearly decrease, pass through zero, and reverse polarity until diode $D2$ is forward-biased. Current will continue through $D2$ until the next power-transfer cycle commences. Note that the behavior is quasi-resonant and that a half-control switch could do the job. However, no such switch will operate at a megahertz rate, which this circuit can reach, so a full-control switch is used, thus suffering no first-order switching losses. A fixed amount of energy is transferred to the load circuit during each energy-transfer cycle, so load voltage and current control is accomplished by the modulation of the transfer-cycle recurrence rate, as in other quasi-resonant designs.

This circuit is another example of a pace-setting concept based on a time-honored process.

19. Thermal Management

Thermal management of a transmitter means controlling its operating temperatures. It also means heat transfer, but this will happen whether there has been proper thermal management or not.

Power will be dissipated throughout the components of a transmitter, but the dissipation will be most significant where energetic electrons strike the inside surfaces of microwave-tube beam collectors or the surfaces of cathode-facing anodes in triodes and tetrodes. Charge carriers impinging on the collectors and drains of transistors are also responsible for their internal power dissipation. The rate-of-arrival of these charge carriers determines current flow, and their kinetic energy, which is proportional to voltage, determines the energy-deposition rate, or power dissipation.

Few electrical, or even mechanical, transmitter designers are completely comfortable with thermal design. Most know only enough "to get burned," so to speak. However, our knowing that literally "getting burned" is a real possibility should heighten our awareness of, if not our ability to solve, all of the thermally related problems.

To begin with, there are only two basic physical mechanisms for heat transfer: radiation and conduction. Radiation, which has its basis in quantum mechanics, involves emission of electromagnetic waves, or photons, from the surface of the hot body. The rate at which heat will be transferred by radiation from a unit of surface area depends upon the difference between the fourth power of the absolute temperature (in kelvins) of the hot-body surface and the fourth power of the absolute temperature of its surroundings, assuming that the area of the surroundings is much greater than that of the hot body. It is also modified by its emissivity, a coefficient anywhere from zero to one that represents the normalized ability to radiate when compared with a perfect "black-box" radiator.

Unlike radiation, which is most efficient in a vacuum, conduction takes place in solids and fluids (both gaseous and liquid) and is related to molecular motion. Heated molecules are more energetic, and they exchange that energy with neighboring molecules that are cooler or less energetic. Convection is a special case of conduction. It can be either "free," or "natural," or forced. It involves the transfer of heat from a solid to a fluid, or from one fluid to another. Convection is also influenced by the ability of a fluid to move freely and to change its characteristics. For instance, the fluid can become less dense when it is heated, or even freeze, liquefy, or vaporize. Also unlike radiation, conduction has its rate of heat transfer figured differently. The rate of heat transfer by conduction for a given unit of area (normal to the direction of heat flow) is proportional to the difference in temperature, ΔT , between source and sink along the path of heat flow. (Although not directly related to electronics at all, the phenomenon of giant tractor-trailer trucks "losing" their brakes has nothing to do with failure of the air system that actuates them but a failure of brake thermal management. If the truck is heavy enough and going fast enough down a steep-enough grade, the heat generated by the brake friction, which is also transferred primarily by radia-

tion, causes the brake temperature to reach a point where the brake diameter will actually increase due to thermal expansion. When the brakes reach the point where they no longer make contact with the drum, the result is another runaway truck.)

19.1 Heat transfer by radiation

Although the primary means of heat transfer in a transmitter is conduction of one form or another, some heat will be transferred by radiation. A classic example of multi-mode heat transfer is the family of glass-envelope, pulse-rated tetrodes that includes the 4PR250 and the 8960, as well as the 4PR125, 4PR400, and 4PR1000 (see Fig. 11-12). These tubes have anodes made of a special metal called Pyrovac ("pyro" for high-temperature and "vac" for vacuum). If you recall, the anodes are cylindrical and are the outermost electrodes. (The physical format is called radial-beam, external anode.) In these tubes the anode is suspended from the top of the glass envelope, where a glass-metal seal provides a connection to external circuitry. For the seal to be properly cooled, a "heat-radiating connector," which is a short cylindrical piece of aluminum with machined fins, must be attached to the external anode connection. If the current terminating on the anode produces enough power dissipation, the lower end of the anode will glow red. The glow starts at the bottom because the thermal-conduction path to the external anode heat dissipator is the longest from that end. As dissipation increases to its rated maximum, the entire outer surface of the anode will glow red. It will literally be red hot. It will now have reached a temperature within the vacuum of the tube at which it can efficiently transfer heat by radiation. (It has been said that there is—or was—at least one tube engineer who could accurately estimate within a few watts the amount of anode dissipation in such a tube just by observing the hue and extent of the anode glow.)

The problem of heat transfer in outer space is also a formidable one because there is no external molecular medium to receive conducted heat. It is possible to build microwave vacuum tubes using collectors that can operate at temperatures sufficiently high to radiate heat directly into space. Temperatures up to 900°C (1173 K) have been tolerated without serious problems. Solid-state devices, however, must be operated with junction temperatures below about 350°C (623 K), and even then their life expectancies can be seriously shortened. The microwave tube collector enjoys a radiation-cooling effectiveness that is almost 24 (or 16) times that of the solid-state device (and even that assumes that the latter's junction itself can be made to directly radiate into space). In the early days of the Strategic Defense Initiative, it was estimated that the surface area of a radiator required to transfer heat from a solid-state transmitter in a space-based radar, with tolerable surface temperature, would equal that of a football field—including the end zones!

Devices producing RF power that are placed on board rockets using liquefied gas as fuel do have a molecular medium with which to exchange heat (and a cold one at that). Unfortunately, transistors become sluggish and virtually inoperative at cryogenic temperatures, except for a special class called the static-induction transistor. However, recent actual experiments have shown that even conven-




Figure 19-1. The elements of thermal resistance.

tional bipolar transistors can be "insulated" from the cold sink in such a way that they will heat themselves up by their own dissipation to a quasi-functional internal temperature.

19.2 Heat transfer by conduction in solids

No matter what makes up the complete thermal path along which heat is transferred, it invariably begins with conduction in a solid. The anode of a triode

Table 19-1. Thermal resistivity of selected materials.

Material	R (°C-in./W)	Material	R (°C-in./W)
Diamond	0.06	Silicon carbide	2.3
Silver	0.10	Steel (300 series)	2.4
Copper	0.11	Nichrome	3.0
Gold	0.13	Carbon	5.7
Aluminum	0.23	Ferrite	6.3
Beryllia ceramic	0.24	Pyroceram (9606)	11.7
Molybdenum	0.27	Epoxy	24
Brass	0.34	Quartz	27.6
Silicon	0.47	Glass (7740)	34.8
Platinum	0.54	Silicon thermal grease	46
Tin	0.60	Water	63
Nickel	0.61	Mica	80
Lead-tin solder	0.78	Polyethylene	120
Lead	1.14	Teflon	190
Boron nitride	1.24	Nylon	190
Alumina ceramic	2.13	Silicone rubber	190
Kovar	2.34	Air	2280

or tetrode or the collector of a microwave electron tube is made of metal in the form of a thin-walled cylinder with a top on it. If that metal is part of the vacuum envelope of the tube, it is usually oxygen-free, high-conductivity copper. Heat generated by the sudden stop of electrons on the inner surface of an anode or collector must be transferred to the outer surface before it can be exchanged with another medium. A well-designed collector and electron-beam system will have power distribution over the inside surface as nearly uniform as possible. The temperature rise (or drop) that enables the heat flow through the thin metal wall will depend upon the total area normal to the direction of heat flow, the length of the path in the direction of heat flow, and the thermal resistivity of the material. If the material is copper, its thermal resistivity is $0.11^{\circ}\text{C}\cdot\text{in.}/\text{W}$.

The thermal resistance of the anode or beam collector will be given as the product of resistivity and path length divided by surface area. Its units are defined as $^{\circ}\text{C}/\text{W}$. In Fig. 19-1, the thermal resistance of the rectangular segment illustrated is $R \times T/A$. Table 19-1 also lists the thermal resistivities of some commonly used materials. Note that the diamond has the lowest resistivity of all. For this reason, the microscopically small junctions of millimeter-wave solid-state devices such as the IMPATT diode are mounted on diamond heat sinks, which are, in turn, swaged into copper blocks. (Note also that the resistivities of the two most common ultimate heat sinks, water and air, are not low at all. In fact, static air is well known as an excellent thermal insulator.)

The relationship between thermal resistivity and thermal resistance is completely analogous to the relationship between electrical resistivity and electrical resistance. In the thermal equivalent of Ohm's law, temperature difference is the "voltage," thermal resistance is electrical resistance, and rate of heat transfer is the "current." The rate at which heat must be transferred, or power dissipated, is usually fixed. To minimize "voltage-drop," or temperature differential, the thermal resistance of the path must be minimized.

The electrical analogy extends to transient phenomena as well. In pulse-modulated transmitters, power is not continuously dissipated in anodes or collectors. Dissipation can occur during very short intervals that periodically recur. But conduction of heat takes time. When power is dissipated over very short periods, adiabatic heating occurs. In other words, temperature rise depends on the mass of the heated material. What the heated mass is—whether it is inner surface of an anode or collector—depends on diffusion depth, which may be very small. What results is nearly instantaneous heating of the surface to a temperature that may be many times as great as the time-averaged temperature. This is followed by an exponential decay of temperature, analogous to an R-C discharge of an electrical circuit, where the discharge exponent is the thermal time-constant rather than the RC product. However, the temperature rise may not be completely adiabatic. The longer the pulse duration, the greater the chance that some heat will be conducted away while energy is being deposited. This condition leads to lower the end-of-pulse temperature rise for the same pulse-energy dissipation.

Solid-state junctions are especially endangered by transient temperature rise because of their intrinsically small thermal mass. The limiting thermal stress for these devices is usually per-pulse energy for operating duty factors at least up to

10%. Average-power dissipation hardly matters at all. This is why practical, solid-state transmitter components operate at as high a duty factor as possible (with 30% being typical) and with pulse durations as long as possible.

19.3 Heat transfer by free, or natural, convection in air

In the vast majority of heat-transfer situations, the ultimate heat sink will be ambient air. It is cheap, abundant, and will move under its own power. Because hot air is less dense than cool air, it rises when it is heated. This proclivity alone makes air effective as a heat-removal medium. (Static air, as mentioned before, is a good thermal insulator.) Natural air convection is almost something for nothing. However, as Fig. 19-2 shows, air will not only rise along a vertical surface, it will be forced away from it to create a boundary layer. The resulting air-velocity profile in the boundary layer shows that the greater the height of the heat source, the farther away from that surface will be the fastest-moving air. The air at the surface remains practically stagnant. The temperature profile, consequently, will show that surface temperature increases the higher up the measurement is taken.

The rate of heat transfer is given as $Q = h A \Delta T$, where Q is the rate of heat





Figure 19-2. Heat transfer by natural convection of air.

transfer in BTU/hour (using the English system of units), h is the convection heat-transfer coefficient in BTU/(ft² × °F), A is the surface area of the hot body in square feet; and ΔT is the temperature difference between the surface and the air in degrees Fahrenheit. Note that h is a coefficient. Sometimes it is called the film conductance or the film coefficient, but it is anything but a constant. It will be different for almost every situation and can vary by at least a factor of 5 for the case of free convection of air. In the case of the vertical hot surface, h varies with both ΔT and L , the vertical height of the hot surface, so that $h = 0.27(\Delta T/L)^{0.25}$. This gives

$$Q = 0.27 \left(\frac{\Delta T}{L} \right)^{0.25} \times A \times \Delta T = 0.27 \frac{\Delta T^{1.25}}{L^{0.25}} \times A .$$

If we know power dissipation, or Q , the temperature rise is given as

$$\Delta T^{1.25} = \left(\frac{Q}{0.27A} \right) \times L^{0.25}$$

$$\Delta T = 2.85 \left(\frac{Q}{A} \right)^{0.8} \times L^{0.2} .$$

Consider the example of the power resistor shown in Fig. 19-2. If it dissipates 20 watts, what will be its average temperature rise? To convert 20 watts into BTU/hour, we multiply it by the conversion factor of 3.41 and get 68.2 BTU/hour (or 34.1 BTU/hour on each side). The area of each side is 10 in. × 2 in., or 20 in.², or 0.14 ft². The height, L , is 2 in., or 0.167 ft. Substituting these values into the equation above, we find that the calculated temperature rise will be

$$\Delta T = 2.85 \left(\frac{34.1}{0.14} \right)^{0.8} \times 0.167^{0.2} = 162^\circ\text{F} .$$

Remember, this is *average* temperature rise. The resistor will be hotter at the top than at the bottom.

19.4 Transfer of heat by forced convection of a cooling fluid

For most high-power transmitter equipment, ultimate heat transfer does not result by natural convection of air or water because of their relatively high thermal resistances. In such applications, the natural tendency of a heated fluid to rise must be augmented by external means such as a water pump or an air blower. This forced flow is also confined within tubelike channels, such as the one shown in Fig. 19-3. As in heat transfer by natural convection of air, the rate of heat transfer Q is still the product of h , the heat-transfer coefficient; A , the surface area of the inside of the tube; and ΔT , the temperature difference between




Figure 19-3. Conditions for heat transfer by forced convection of a fluid medium.

the tube's inner surface and that of the fluid. The problem is that the value of h is subject to even more different variables of the system, including tube diameter and fluid velocity, viscosity, specific heat, and thermal conductivity. In addition, not only will the fluid have a higher temperature when it exits the heated length (T_2) than when it enters (T_1), the tube will have a higher temperature at the exit end than at the entrance. There will be not only a T_1 and a T_2 but a ΔT_1 and a ΔT_2 as well. This is handled mathematically by invoking a log-mean temperature difference,

$$\Delta T_M = \frac{\Delta T_1 - \Delta T_2}{\ln \left(\frac{\Delta T_1}{\Delta T_2} \right)}$$

The value of h depends on the velocity profile of the fluid in the radial direction and the degree of fluid turbulence. The last thing we want is maximum velocity down the center of the tube and stagnant fluid at the surface, the so-called Holland-Tunnel effect. The degree of turbulence in the fluid flow—in other words, a more uniform radial velocity profile—is quantified by the Reynolds number, N_{Re} , which is defined as the product of the tube diameter, D , and the fluid velocity, V , divided by the fluid viscosity, u . The value of h can be extracted from the dimensionless Nusselt number, Nu_m , which is defined as the product of h and tube diameter, D , divided by the thermal conductivity, k . The value of Nu_m depends on the Reynolds number, N_{Re} , and the Prandtl number, Pr_m , which is equal to the product of the specific heat, C_p , and viscosity, u , divided by the thermal conductivity, k .

Figure 19-4 shows how the Nusselt number varies with Reynolds number. The non-abrupt dividing line between laminar and turbulent flow occurs at Reynolds number 2300. For N_{Re} below 2300, Nu_m varies as




Figure 19-4. How the Nusselt number, which includes the heat-transfer coefficient h , varies with the Reynolds number.

$$\sqrt[3]{N_{Re} \times Pr_m \times \frac{D}{L}}$$

For N_{Re} above 2300, Nu_m varies as $N_{Re}^{0.8} \times Pr_m^{0.4}$. As can be easily understood, quantitative analysis of cooling effectiveness depends almost entirely on the determination of h , which in forced-fluid convection can vary by a factor of 40 or so. In general, however, the relative weights of the fluid factors on the heat-transfer coefficient can be expressed as

$$h = 0.23 \times \text{density}^{0.8} \times \text{thermal conductivity}^{0.6} \\ \times \left(\frac{\text{specific heat}}{\text{viscosity}} \right)^{0.4} \times \frac{\text{fluid velocity}^{0.8}}{\text{fluid - channel diameter}^{0.2}}$$

19.5. Closed-loop water-cooling systems

Most transmitter engineers are not going to be involved in designs that relate to what has been written so far about thermal management. The devices that must be cooled are already identified, and their cooling requirements are usually specified in terms of air-flow rate as a function of altitude or water-flow rate and water-course pressure drop. Also specified will be the allowable maximum fluid temperature at the inlet—and sometimes minimum inlet temperatures as well. If these external specifications are met, we can be reasonably assured that internal operating temperatures will be in the proper range.

In a forced-flow liquid-cooling system, there is an important quantitative rela-




Figure 19-5. Performance of water-cooling system in heat removal.

tionship that most of us are aware of. It is that the heat-transfer rate Q equals the flow rate of the liquid times the density of liquid times the specific heat of liquid times the difference between the outlet liquid temperature and the inlet liquid temperature. For water, the expression reduces to the familiar heat-removal rate, which is expressed in watts and is equal to 264 times the flow rate in gallons per minute times ΔT in degrees Celsius. This is shown in Fig. 19-5 in nomograph form for water and for various concentrations of water and ethylene glycol (anti-freeze). This is the basis for estimating anode or collector dissipation by measuring inlet and outlet temperature and flow-rate of the cooling water. In calorimetric RF "dummy" loads, this method is used to measure average input power, which is often transmitter-output power as well.

If air that has been forced through a tube is the cooling fluid, the relationships are more complicated. What matters is the mass of the air that passes through all of the cooling tubes, or channels, of the air-cooled device per unit of time. Typically this is expressed in pounds per minute. Unlike water, however, air is compressible, and its density varies with temperature and barometric pressure. Moreover, what is usually measured is not mass flow, but volumetric flow, such as gal./min for water or ft³/min for air. The heat-transfer rate for air, which is comparable to the expression given above for water, is

$$Q = \text{watts} = \frac{164 \times \text{flow rate} \left(\frac{\text{ft}^3}{\text{min}} \right) \times \Delta T (^\circ\text{C})}{\text{inlet temperature (K)}}$$

at 1 atmosphere (29.92 in. of mercury). When the expression is evaluated at, say, an air-inlet temperature of 55°C, the absolute temperature is 328 K and the heat-transfer equation reduces to $Q = \text{watts} = 0.5 \times \text{flow-rate (ft}^3/\text{min)} \times \Delta T (^\circ\text{C})$. As inlet air temperature increases, air density decreases and mass flow decreases as well. If the altitude is higher than sea level, the density must be adjusted for that factor too. At any temperature or pressure, air density (expressed in lb/ft³) will be 0.737 times atmospheric pressure (in inches of mercury) divided by absolute temperature (in kelvins).

Most air-cooled vacuum tubes have an upper temperature limit of 250°C. It is usually measured at the core of the "radiator" (more properly named the "forced-convectator"), which is attached to the anode or collector of the device. Characteristic curves are usually supplied by the manufacturer indicating the required air-volume flows in ft³/min for varying anode power dissipation. Also, in order that an appropriate air blower can be specified, the tube maker includes pressure-drop data through the cooling channels as a function of flow rate.

What neither of the relationships given above suggests, however, is how hot the inside surfaces of the anode or collector—or the dummy load, for that matter—will get. The ΔT used in this case is parallel to the cooling flow, or end to end. In Fig. 19-3, if L represented the entire length of the cooling path, ΔT would be $T_2 - T_1$. What is not revealed is the ΔT perpendicular to the cooling flow, which represents the temperature rise from fluid boundary to heated surface,




Figure 19-6. Components of a typical closed-loop, high-purity water-cooling system.

part of which is the log-mean ΔT described earlier. The total temperature rise is determined by the total thermal resistance of the heat-flow path, which is essentially transverse to the flow of liquid coolant.

A closed-loop water-cooling system for a high-power transmitter resembles in many ways the cooling system for a car engine. The major components of such a system are shown in Fig. 19-6. One of the key components is the pump, which provides the pressure head that enables the forced flow of coolant. (Although the liquid coolant could circulate through the system by natural convection from low point to high point and back, it will do so at too leisurely a pace.) Other key components include the heat exchanger, which accomplishes the second and ultimate level of convective heat transfer from liquid to ambient air, and the blower, which provides the forced flow of air through the heat-exchanger coils. Where there is an abundant source of ambient water, such as on board a ship or near a coastline or lake, water-to-water heat exchangers are often used. These are capable of heat transfer at lower ΔT and usually to a heat sink (the ambient water) that has a lower temperature than that of the ambient air. The disadvantage of the water-to-water exchanger is the very real possibility of leaks between the two water paths. A mandatory requirement for the water (or alternative liquid coolant) in the closed-loop primary path is that it be high-purity, de-ionized, and oxygen-free. Any leak from the secondary side of the exchanger, even if the larger pieces of flotsam have been filtered out, will make short work of polluting the primary side.


In the interest of maintaining the purity of the primary-loop coolant, a purity-maintenance loop usually shunts the main flow path. Included in its flow path are mineral-bed purity cartridges and submicron particle filters. The purity of the water is usually monitored by measuring its electrical resistivity. This function is carried out by the resistivity meter, which is shown shunting the conductivity cells. High-purity water has a resistivity typically greater than 1 megohm-cm, although marginal performance is obtained with resistivity as low as 300 kilohm-cm.

The cooling paths of most microwave-tube transmitters do not include loads that may operate at high voltages with respect to ground. The anodes of large hard-tube-modulator switches may undergo large pulsed-voltage swings, however; and the anodes of electronic voltage-regulator series tubes will also be at high voltage with respect to ground. In these cases not only is water resistivity important, but the length and diameter of the water columns leading to and from the water-cooling jackets are important too. The combination of the three will determine the electrical resistance of the water path and the amount of electrical current that will flow through it. Such current can produce electrolysis effects in the water. Needless to say, the water conduits must also be non-conducting. Standard polyvinyl chloride (PVC) pipe or, to be safe, chlorinated PVC (CPVC) pipe is very useful in this application. Where high flow rate dictates large-diameter pipe, the length of the water column must be proportionately long. Before PVC pipe became a useful alternative, and where high voltage was also involved, porcelain-ceramic water-columns in the shape of cylindrical coils were almost standard water-system components. Today, square-sided coils are often fashioned from short straight sections of PVC pipe and 90° plumbing elbows, thus allowing a long flow path to be coiled up in a relatively compact volume.


The materials used in a high-purity water loop should be restricted to copper, bronze, brass, stainless steel, and plastic, like the PVC pipe. No ferromagnetic materials should ever be used. The other important feature, and one not likely to be found in an automobile cooling system, is the nitrogen-gas blanket, which is intended to replace the air in the unfilled volume at the top of the liquid reservoir. The remaining components are self-explanatory: interlock switches, temperature and flow gauges, flow-adjusting valves and shut-off valves, reservoir liquid-level indicator, and pressure regulators and gauges.

The one condition throughout the cooling system that needs to be interlocked with the electrical control of the transmitter is water flow. Several kinds of different flow meters and flow-interlock electrical switches are available for this purpose, as illustrated in Fig. 19-7. The prototype of all of them is the classic variable-area flow meter illustrated in Fig. 19-7a. Called a Rotameter, it uses a float that is suspended in a conical-shaped glass pipe by the flow of liquid. The greater the flow, the higher up the pipe the float will be suspended. A feature on the float allows it to be compared with a scale, which may be either external or etched on the glass. Needless to say, the Rotameter only works when it is vertically oriented.


A spring-loaded version of the Rotameter (Fig. 19-7b) uses constant-diameter pipe, but its metering cone is tapered with its diameter decreasing along the direction of flow. Liquid, entering through holes in the slug that holds the meter-




A. Classical variable-area flow meter (Rotameter); must be vertically oriented




B. Practical spring-loaded Rotameter (orientation insensitive)



C. Piston-type variable-area flow meter and flow switch



D. Vane type variable-area flow meter



E. Turbine flow meter

Figure 19-7. Types of liquid flow meters and flow-interlock switches.

ing cone, applies pressure on an annular piston whose motion is opposed by a calibrated spring. The greater the flow, the greater the area required between the inside of the piston and the outside of the metering cone for positional equilibrium, meaning that the piston must move along the direction of flow. The movement of the piston can be magnetically coupled to an external annular follower whose position can be related to flow by a longitudinal scale. The magnetically coupled follower can also be used to actuate a small electrical switch, whose longitudinal position can be adjusted to give an electrical indication of whether flow is above or below a preset threshold.

Another piston-type, spring-loaded, variable-area flow meter (Fig. 19-7c) can be configured to directly drive a meter needle and/or an electrical interlock switch. The vane-type variable-area flow meter (Fig. 19-7d) uses a semi-circular vane in a not-quite hemispherical chamber to achieve variable-area performance against a flat-spiral spring. The resultant deflection is rotational, which can be compared with a calibrated meter scale and operate an electrical interlock switch through a cam, as shown.

The turbine-type flow meter (Fig. 19-7e) is probably the most inherently accurate, but it is also the most expensive. In it, liquid flow causes a turbine to rotate, as flowing gas does in a jet engine. The turbine's angular velocity is indicated by a magnetically coupled response to the passage of each rotor-blade tip past a magnetic pickup, giving a variable-frequency pulse-train electrical output. The frequency, or course, is proportional to flow rate.

Such flow meters and interlock switches are usually deployed in the return lines of each coolant path. (It does not matter how much liquid starts out in a cooling path, only how much comes back.) A typical water manifold for a dual-klystron RF power source is shown in Fig. 19-8. The main water-flow path, whose flow rate is 47.5 gal./min, is directed to the collectors of the two klystrons. Water flow is connected to the collectors in cascade, or series, because of their relatively low individual pressure drops. A separate path cools the body of each klystron because of relatively high pressure drop at the 2.6 gal./min flow rate. A third parallel path, also 2.6 gal./min, cools the focus electromagnets of the two klystrons. They too are connected in cascade. The four return paths each have separate temperature gauges, electrical temperature-interlock switches, integrated flow meters, and interlock switches. Total supply pressure and return pressure are also metered. Check valves in the lower flow paths prevent backflow when those water courses are drained.

19.6 Vapor-phase cooling, or convection with state change

Water boils at 100°C. In so doing, it changes state from a liquid to a vapor called steam. To accomplish this change of state, even at constant 100°C, it must absorb a considerable amount of energy known as the latent heat of vaporization, which is 540 joules per gram of water. While water is in the liquid state, it takes one joule of energy to raise its temperature one degree Celsius. The usual outlet-inlet temperature differential for a closed-loop, forced-convection water-cooling system is in the range of 30°C to 40°C. A temperature of 70°C is typically considered a maximum outlet temperature because it is low enough to preclude spot-boiling within a water course, which can reduce the heat-transfer effective-

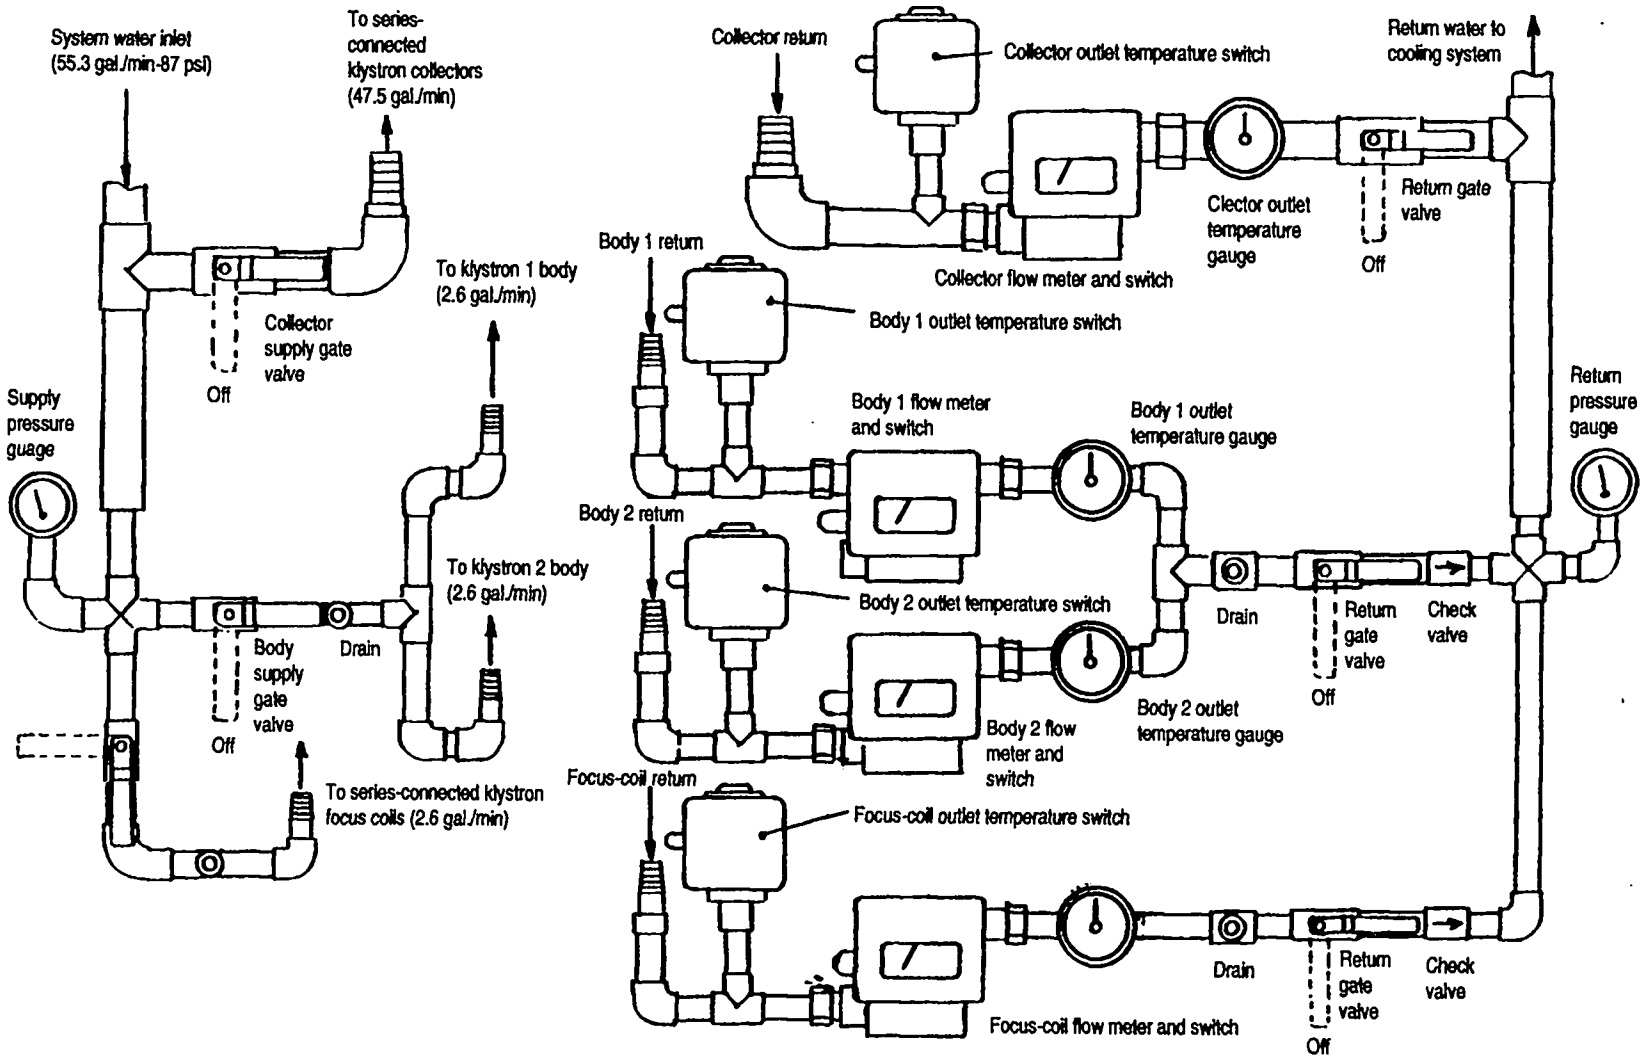


Figure 19-8. Typical water manifold arrangement for a dual-klystron amplifier system.




Figure 19-9. Components of a vapor-phase cooling system for high-power vacuum tubes.

ness at the surface, especially if the flow rate is marginal. This means that each gram of water is responsible for removing between 30 and 40 joules of energy, compared with 540 joules removed for each gram of water if it is allowed to be converted into steam. Energy is not power, however, and we are interested in the *rate* at which energy can be removed, which is characterized in terms of power density. By forcing water through the cooling channels, high mass flow can be obtained (as well as the desired attributes of turbulent flow). If each gram of water can take away 30 to 40 joules of energy, and we move a lot of grams past a given point per unit of time, the power density can be quite great, typically 1000 W/cm^2 of internal anode (or collector) surface. By contrast, the power density of a closed-loop vapor-phase cooling system, which is powered by natural convection, can be at maximum 135 W/cm^2 of internal anode surface.

This is still pretty good, considering that a vapor-phase system, such as the one shown in Fig. 19-9, can be much simpler than a closed-loop, pump-driven water-cooling system. And it uses a lot less water. Primary heat transfer takes place in the boiler, into which the anode or beam collector of a tube is inserted and sealed so that the system will be vapor-tight. The tube's anode and boiler axes must be vertical, and the water level must be as shown. The function of the control box, which acts as a manometer in that the water level in it should be the same as that in the boiler, is to maintain the proper water level. It does so by opening the solenoid-operated valve to let make-up water into the control box from the reservoir as needed. If this does not do the job, or does not do it rapidly enough, the low-water alarm and the electrical interlock will signal the transmitter power source to shut down. An equalizer line, connected between the boiler steam outlet pipe and a comparable outlet port on the control box, makes sure that the steam vapor pressure is the same at the top of the control box as it is at the top of the boiler. The steam gives up the latent heat of vaporization that it




Figure 19-10. Temperature characteristics of vapor-phase cooling system.

has accumulated to either ambient air or secondary cooling water in the condenser. Then it reverts to the liquid state and rejoins the system at the water-inlet side to repeat the process.

The graphs in Fig. 19-10 show the same information, but in linear coordinates on the bottom and in logarithmic on the top. What both illustrate is the critical temperature dependence of the heat-transfer power-density performance. The maximum of 135 W/cm² is achieved at the critical temperature of about 125°C. The temperature condition at this point or just below it is called nucleate boiling, which is characterized by individual bubbles of steam that break away from the heated surface and rise through the surrounding water, carrying their heat of vaporization with them. If the temperature is allowed to get higher, a phenomenon called calefaction takes place. Here, individual bubbles merge increasingly into a sheath of steam at the heated surface. These bubbles begin to thermally insulate the surface from the surrounding water. The heat-transfer effectiveness steadily deteriorates with rising temperature until the Leidenfrost point is reached.

Here, the vapor sheath is complete, and all heat transfer takes place through it. If the power-dissipation density in the heated surface is not reduced from its 135 W/cm^2 value, the temperature at the surface will rise to 1000°C , which obviously must be avoided.

Avoidance strategies employed in the design of the outside surfaces of vapor-cooled anodes or beam collectors tend to call for thick radial fins. The radial dimension of these fins permits a radial temperature gradient as well. If this gradient straddles the 125°C optimum point—hotter at the root and cooler at the tip—the average heat-transfer performance will be less than the maximum achievable if temperature were everywhere constant, but it will be self-correcting. A hot spot at a fin root will push the heat-transfer performance there further away from the maximum point, but will pull the temperature at the fin tip up, pulling its performance closer to the maximum. Needless to say, the external area of the heated surface will be greater than the internal.

The condenser in the vapor-phase system is more efficient in secondary heat transfer than the heat exchanger in a water-cooled system because of the more advantageous temperature difference. Steam and steam that has condensed back to water are both at 100°C . If the heat is transferred to ambient air at 25°C , ΔT is 75°C . A water-cooling system with a 30°C water-temperature rise and a maximum outlet temperature of 70°C has an inlet temperature of 40°C . The average water temperature in the heat exchanger will be 55°C , giving a ΔT of 30°C with ambient air at 25°C . The steam-water condenser will have something like 2.5 times the temperature-difference of the water heat exchanger. Furthermore, it can be built smaller, lighter, and cheaper. And because water flow in a vapor-phase system is accomplished automatically by what is called thermosyphoning, these systems operate with little or no acoustic noise, which can sometimes be important. Pump-driven, water-cooled systems can never be called quiet. Nor can blower-driven air-cooling systems (especially if they operate at 400 Hz primary power).

19.7 Multi-phase cooling

The most modern high-power tubes are now being cooled with hybrid systems that use a combination of vapor-phase and high-flow-rate water cooling. These systems are generically called multi-phase cooling. A trade name for such a system is Hypervapotron. Multi-phase cooling is characterized by surface vaporization of water into steam bubbles that are then almost immediately condensed back to water by the rapidly flowing surrounding water. The external cooling characteristics of this system are largely indistinguishable from conventional water-cooled tubes, except that its pressure drop is likely to be considerably less than a conventional system for comparable flow rate.

20. Transmitter Control, Monitoring, and Interlocking

No matter how internally complex or how high-power a transmitter might be, its external or human-operator interface might be quite simple—as simple as turn it on, turn it off, or, if the transmitter stops working, to say what’s wrong. The first two controls represent the operator or larger system communicating with the transmitter. The third is the transmitter communicating with the operator or larger system. In addition to providing a way for an operator to turn the system on and off, all transmitters must be capable of monitoring their internal status and automatically taking the proper action in case of anomalies. Such a capability is called a control circuit. In terms of computer logic, a transmitter control circuit will be one form or another of an AND gate. Often it will have literally hundreds of independent inputs, all of which must prove “true” before ultimate transmitter output will be obtained. We begin with the simplest control circuit of all.

20.1 The relay-logic ladder-network control circuit

The great preponderance of existing transmitter control circuits—and even some that are still on the drawing board (or CAD workstation)—are nothing more than expanded versions of the bare-bones illustrative circuit shown in Fig. 20-1. This is the classic “daisy-chain” ladder-logic-type AND gate with indicator lamps at every rung. As shown, it is capable of automatically sequencing transmitter turn-on from the very beginning to the point where the transmitter is ready for high-voltage dc to be applied to the transmitter output device. This is basically accomplished by applying control power at the beginning of the chain. (The circuit logic also assumes that low-voltage primary power is continuously connected to it.) Let’s follow the daisy chain.

If the transmitter is water-cooled, a liquid-level switch in the water reservoir will be closed if the level is adequate. The switch closure applies power to the first indicator lamp (green) and to a power contactor that applies primary power to the cooling-system pump, which we do not want to turn on if there is not enough water in the system. Simultaneously, it powers the next rung in the ladder. (Any cooling-system air blower has no preconditions that need to be satisfied; it can be turned on coincident with control power.)

Assuming that the RF device is magnetically focused and that the focus coil is water-cooled, the next event of importance is water flow in the coil. If it is adequate, the water-flow interlock switch will be closed, applying control power to the next indicator lamp, to the coil of a contactor that applies primary power to the focus-coil dc power supply, and to the next rung.

If the RF device has a focus coil, it usually has a body and a collector that also need to be water-cooled. The next switches in line are the body- and collector-water-flow interlock switches, which are closed if flow is adequate. If these switches are closed, power is applied to their respective indicator lamps, to the contactor that applies filament power to the RF device, and to the next rung.

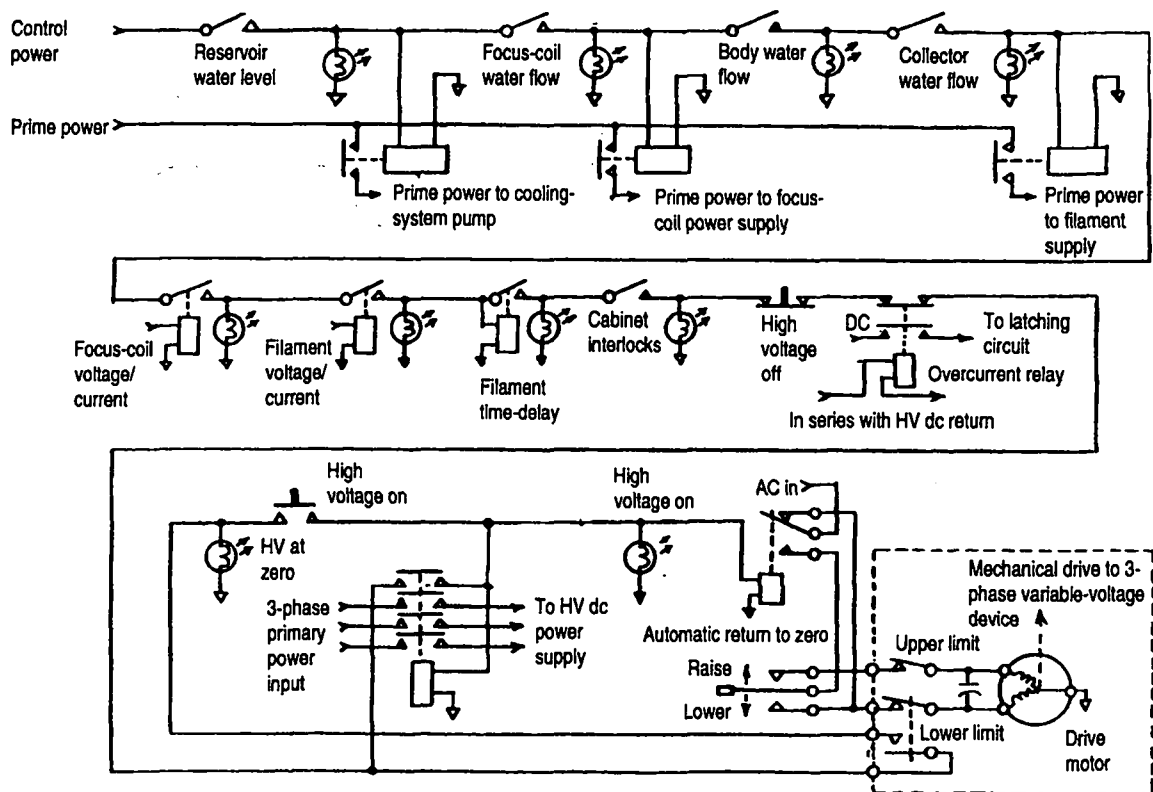


Figure 20-1. A simple series-string, or daisy-chain, type of transmitter control AND circuit.

The next two rungs are closed by relay contacts. The first is closed when its relay is energized with adequate measured current and focus-coil voltage. (In today's world of current-regulating dc power supplies, it is not enough to measure current alone, because even a short-circuited coil could draw the proper amount of current.) The second is closed by a relay that is energized by adequate filament voltage and current.

When both filament voltage and current are at their proper values, a cathode-heating time-delay relay is energized. For large-area, indirectly heated cathodes, this heating time can be as much as 15 minutes. Once the delay has elapsed, the relay contacts close, lighting the "delay-complete" indicator and sending control power to the next rung, which is the last fully automatic one in our simplified control ladder. Nevertheless, it is by far the most important. It is shown as a cabinet-interlock switch, but it represents all of the personnel-safety interlocks and interrupts that are not shown—including the kinds of switches described in Chapter 4, Section 6—and any mushroom-button emergency-off switches. Once this daisy chain has been activated, the cabinet-interlock indicator lamp is illuminated and the high-voltage on/off operator commands are enabled.

The high-voltage on/off controls illustrated are of the momentary-actuation type. The high-voltage off push button is normally closed and is in series with the normally closed contacts of a high-speed electro-mechanical overcurrent relay, whose actuating coil is in series with the high-voltage dc return conductor. The relay contacts will open if the high-voltage dc return current exceeds a preset threshold. Coupled to the normally closed contacts is a pair of normally open

contacts in series with a latching circuit, which gives a more permanent indication that the overcurrent relay was momentarily tripped. The control path then leads through normally open contacts that are mechanically linked to the normally closed contacts of the lower-limit switch for the motor, which drives the variable-voltage device in series with the input to the high-voltage dc power supply. When the ac input to the high-voltage dc supply is at a minimum, these contacts will be closed, and the "high-voltage-at-zero" indicator (usually amber) will be on.

The closing of these contacts signals the operator that the on push button for the high voltage has been enabled. When it is momentarily depressed, control power will be applied to the three-phase contactor that connects primary power to the high-voltage dc power supply. An auxiliary set of contacts (the uppermost ones) short-circuits both the high-voltage-on push button and the high-voltage-at-zero contacts, latching the three-phase contactor in the closed position. The operator may now raise the output of the variable-voltage device, increasing the high-voltage dc applied to the transmitter output device until the desired operating conditions are reached. If nothing else happens between this event and the time when RF power is no longer needed and the high-voltage power supply can be shut off, the operator pushes the high-voltage-off push button, momentarily interrupting the coil power of the three-phase contactor, which opens the auxiliary contacts. Neither the high-voltage-on nor the high-voltage-at-zero contact is still closed, so the three-phase contactor remains open. If an automatic return-to-zero relay is used, as shown, it too will be de-energized, and its normally closed contacts will perform the same drive-motor function as the high-voltage "lower" switch position.

Suppose, however, that in the midst of normal operation an internal malfunction occurs, such as a leak in the water reservoir that drains all of the cooling water. On its way to empty, the water level will pass by the minimum amount required to keep the water-level switch closed. When its contacts open, the effect on the three-phase contactor is the same as pushing the high-voltage-off push button. The high-voltage dc will be dumped. The unambiguous indication given by the interlock ladder is that there is insufficient water level in the reservoir because the water-level indicator lamp is no longer illuminated. Unfortunately, none of the downstream indicators lights is on either. Here lies the basic shortcoming of the daisy-chain interlock string. Its indications are often ambiguous because it suggests that everything else in the circuit has failed too. (A simple burned-out bulb can give a false failure indication as well.) Until the leak is fixed and water-level is restored, there is no way of knowing if anything else has gone wrong. The opening of any interlock contact will dump the high-voltage dc and extinguish its indicator and all of the others downstream of it. (It is only fair to relate that large systems using this type of control sometimes make available to the operator "cheat" switches that permit the bypassing of an interlock switch to see how many of the downstream indications are still functioning. Although the cheat switch is supposed to be used only temporarily, it subverts the integrity of an otherwise fail-safe system. With respect to personnel-safety interlock systems, use of the cheat switch can also raise the issue of gross negligence.) This shortcoming, along with the availability of robust and affordable

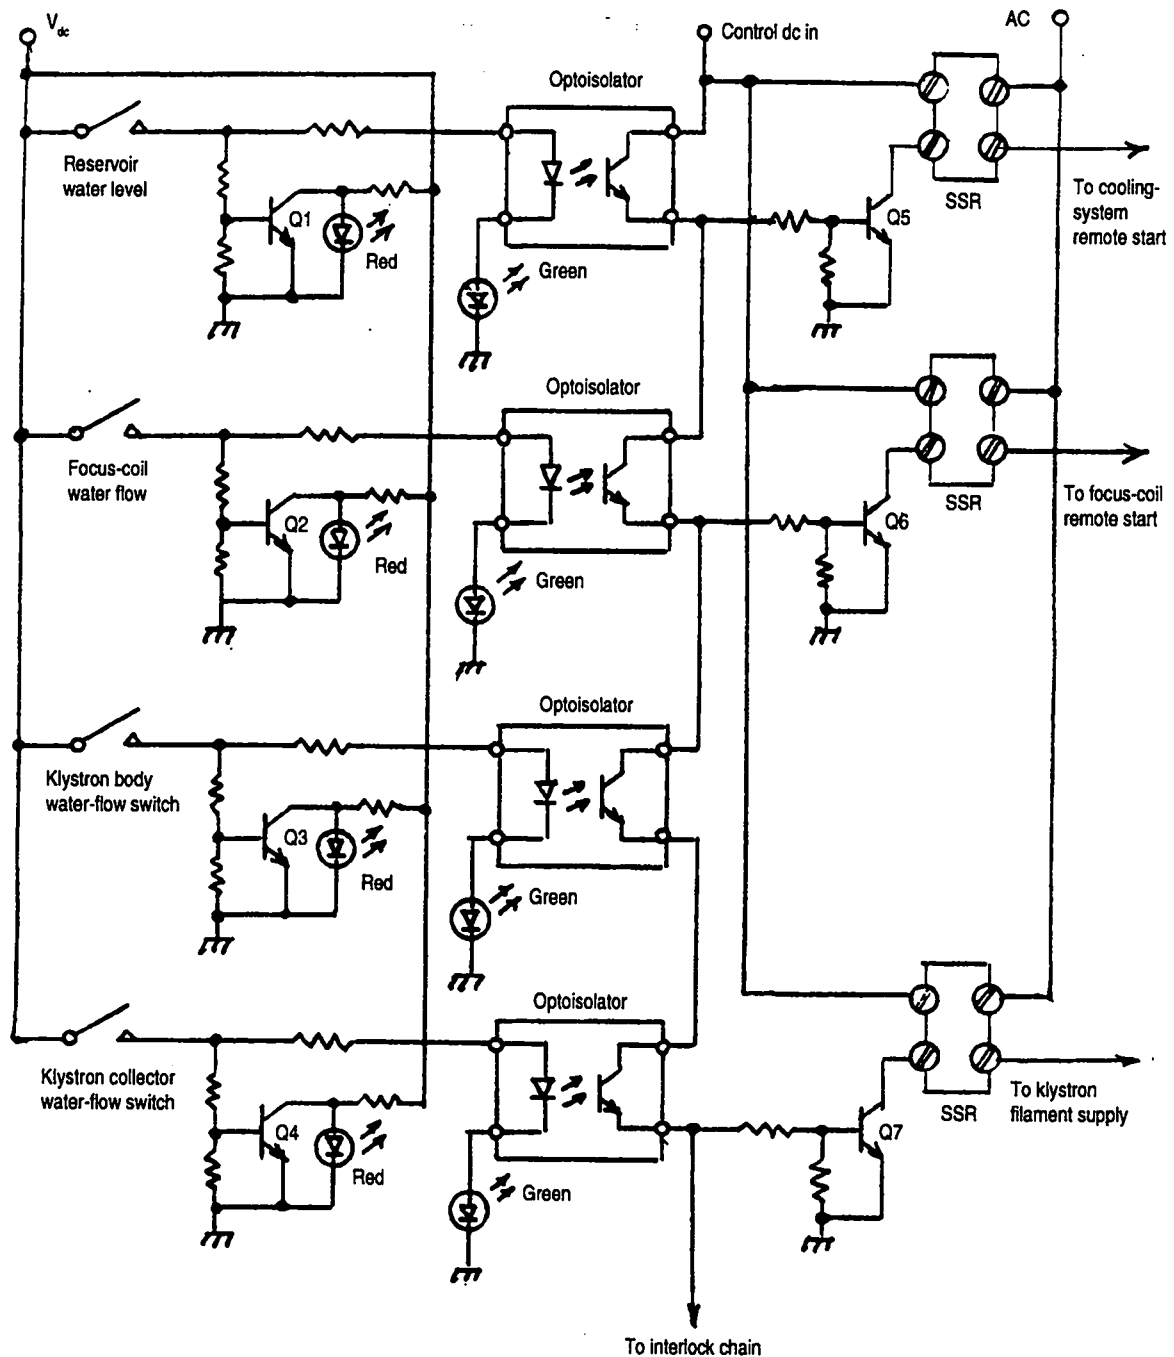


Figure 20-2. A solid-state transmitter-control circuit with series control but parallel status read-out (continued in Fig. 20-3).

solid-state components, has engendered advanced control-system topologies.

20.2 A solid-state, serial-control, parallel-indication interlock circuit

Figures 20-2, 20-3, and 20-4 show some of the features of a hard-wired, dedicated-circuit, solid-state transmitter-control system. It features a series-connected, ladder-network AND gate like that of the daisy chain and full-time status indication for all of the monitored functions. It uses dual indicators for each function—

green for "true" and red for "false"—to mitigate the burned-out-bulb ambiguity. If there were only green "true" indicators and one of them burned out while the interlock chain sequence was in some intermediate state, the light-out indication could be interpreted as either failure of the function or failure of the indicator. This is not the case with dual indicators. If neither the red nor green indicator is illuminated, the condition most likely suggests that one indicator or the other (or, more unlikely, both) is faulty, but not the function. Whatever ambiguity remains for this condition is lessened even further in the true daisy chain because if the downstream indicators are still on, it is most likely the indicator that has failed. To a large extent, the light-emitting-diode (LED) indicator has made the burnt-out-bulb syndrome a moot point because the LED tends to last as long as almost any other component in the system. On the other hand, incandescent, 24-Vdc bulbs, which are found in many older control systems, require almost continual replacement.

The building-block circuits are first shown in Fig. 20-2. The basic transducers and interlock switches are virtually the same as the ones used in the previous topology, starting with the water-reservoir level switch. Instead of being part of a series-connected string, however, its contacts are associated with an individual solid-state circuit. When reservoir level is too low, the contacts are open and control voltage V_{dc} produces current flow through the first red LED. When the contacts close, V_{dc} produces base current in $Q1$, whose collector short-circuits the red LED, thus turning it off. It also produces current through the photodiode of the first optically coupled isolator and the first green LED, which are connected in series. Note that the sum of these currents is not very large (in the milliamperage range). The switches and sensors are operated in the dry-contact mode, so failure through contact welding is a virtual impossibility. (Failure of the switch to conduct when it is closed is the more likely failure mode, but it is fail-safe.)

The photo-sensitive transistor of the optoisolator is now turned on, allowing current from the control dc input to the base of $Q5$, turning on the first SSR and also enabling current flow to the next optoisolator phototransistor just below it. The turning on of the dc control input to the first SSR through $Q5$ switches ac to the cooling system remote-start function.

The subsequent functions are the same as those shown in the daisy-chain system. The AND gate is formed by the series connection of all of the optoisolator phototransistors. Subordinate actions take place with the same prerequisites as before. Transistor $Q6$ will not enable the focus-coil power supply SSR until both water-level and focus-coil water flow are "true." Transistor $Q7$ will not enable the klystron filament supply SSR until klystron-collector water flow and all functions before it are "true."

The optoisolator chain continues on Fig. 20-3, but an important high-speed interrupt has been added to the sequence that responds to excessive klystron body current. The klystron illustrated has a collector that is electrically isolated from the body (without which there is no way to measure body current). The positive (low-voltage) return lead of the high-voltage dc power supply is connected to the collector through current meter $M3$, which indicates average collector current. The return is also connected to the klystron body (ground) through current meter $M4$ and a series resistor. Any ground or klystron-body current

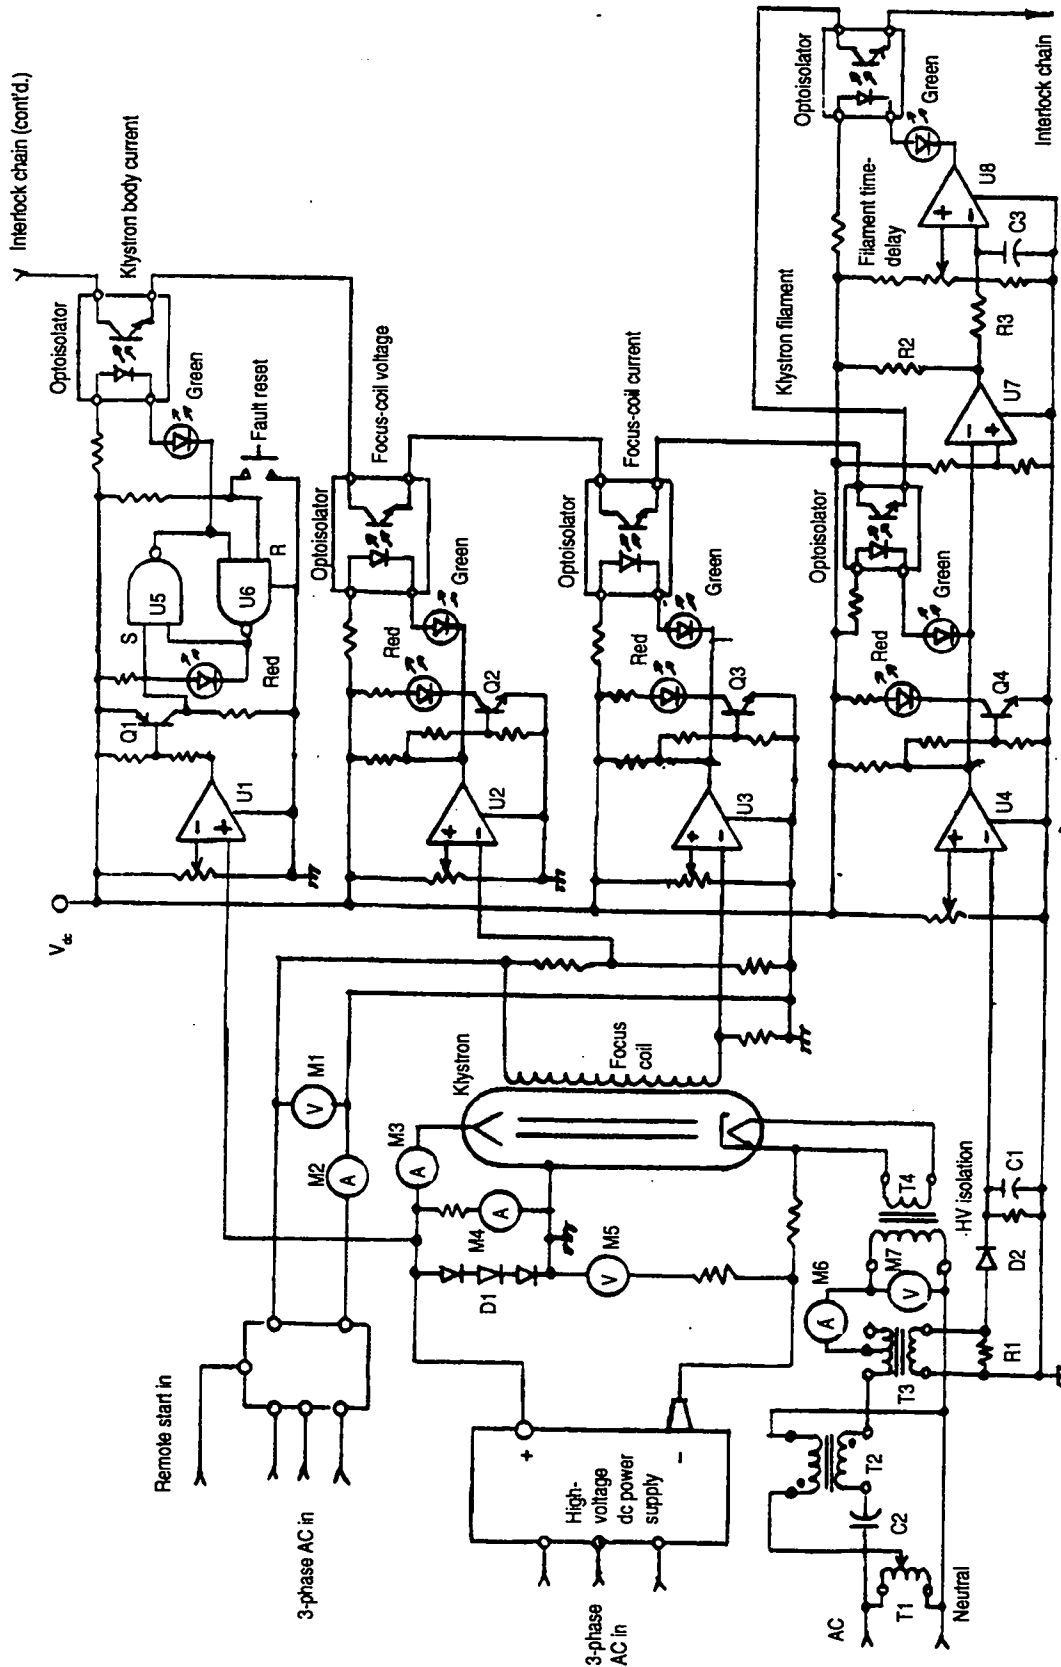


Figure 20-3. Continuation of solid-state, serial/parallel control circuit (continued in Fig. 20-4).

must return to the power supply through $M4$, which indicates its average value. The voltage across $M4$ and the resistor in series with it is also proportional to instantaneous body current. This voltage is applied to the positive input of an integrated-circuit voltage comparator, $U1$ (such as the ubiquitous LM339). An adjustable-threshold voltage is applied to the negative input. As long as the klystron-body current is below an acceptable upper limit, the negative input of the comparator will be more positive than the positive input and the output of the comparator will be low, meaning that its open-collector output device will be conducting. This produces base current in the PNP transistor $Q1$, turning it on and pulling the voltage at its collector positive to near the V_{dc} rail. This voltage is also connected to the "set" (S) input of the "set/reset" flip-flop (bi-stable multivibrator), which is created by interconnecting the two integrated-circuit, dual-input NAND gates, $U5$ and $U6$, as shown. When both input gates of the NAND gate are high, the output will be low. If either input is low, the output will be high. In the reset (R), or normal, state of the flip-flop, the output of $U6$ is high, which is one input to $U5$. The other input is the collector voltage of $Q1$, which is also high. So the output of $U5$ is low, assuring that the output of $U6$ will be high and permitting current through both the optoisolator photodiode and the green LED indicator.

Should the klystron-body current even momentarily exceed a value that makes the voltage at the positive input of $U1$ exceed the threshold voltage at the negative input, the output of $U1$ will go high, $Q1$ will stop conducting, the S input of $U5$ will go low, the output of $U5$ will go high, current will stop in the photodiode and the green LED, and the interlock chain will be opened with microsecond speed. In addition, the output of $U6$ will go low because both of its gate inputs are now high, turning on the red LED, which indicates a fault in this channel. The circuit is reset by pressing the "fault-reset" button, which momentarily grounds the R input to the flip-flop, making the output of $U6$ high again. The klystron body current will be zero at this time because the high-voltage dc has been shut down, so this will reset the original conditions.

If the klystron-body-current overload was caused by a high-voltage arc in the klystron electron gun, all of the arc current will be body current, and it can have a peak current that is the operating beam voltage divided by the value of surge resistance in series with the cathode. The klystron illustrated has continuous beam current (it is not pulse-modulated), and might have a beam voltage of 30 kV or so. If the cathode surge resistor is 30 ohms—a value picked for an obvious reason, but it is by no means atypical—the peak instantaneous arc current could be 1000 A. This is the reason for the series diode string shown as $D1$, which shunts $M4$ and its series resistor. The forward voltage drop of the diode string, even at low current, is chosen to be greater than the threshold voltage of $U1$. Diode $D1$ is expected to clamp the voltage across it to a value safe for the positive input of $U1$ while diverting nearly all of the 1000 A of fault current through itself.

The next two rungs in the optoisolator ladder are focus-coil voltage and current, both of which use identical electronic circuits. A sample of the focus-coil voltage that is taken from a voltage divider is applied to the negative input of comparator $U2$, and a voltage proportional to focus-coil current, which is obtained from a sampling resistor in series with the focus coil, is applied to the

negative input of comparator $U3$. The positive inputs of both comparators are connected to adjustable threshold voltages. Note that both terminals of the focus coil float with respect to ground. If both the voltage and current samples exceed the threshold voltages, which represent acceptable minima rather than maxima, the outputs of both $U2$ and $U3$ will be low, enabling current flow in their respective optoisolator photodiodes (and their green indicator LEDs) and completing the conduction path through the series-connected optoisolator phototransistors. Should either sample signal fall below its threshold, its comparator output will go high, either $Q2$ or $Q3$ will be turned on (thus illuminating one or the other red "fault" LED), and conduction in the phototransistor chain will be stopped. Meters $M1$ and $M2$ indicate focus-coil voltage and current and are usually accessories of the focus-coil power supply.

The last two rungs shown in Fig. 20-3 monitor klystron-filament current and cathode-heating time-delay. The actual filament current flows in the secondary of transformer $T4$, one side of which is connected to the klystron cathode. The cathode is normally 30 kV negative with respect to ground, thus making it awkward to directly monitor the filament current. (A useful, although not entirely accurate, indication of filament current can be obtained by monitoring the $T4$ primary current, which is usually smaller by the $T4$ voltage step-down ratio and referenced closer to ground. This method is not entirely accurate because of $T4$ magnetizing current.) This is the reason for the presence of $T3$, which is often a low-voltage filament transformer that is connected with its low-voltage secondary in series with the primary current of $T3$. The secondary is center-tapped as shown to provide half the voltage in order to maximize the step-down ratio. A resistor, $R1$, is connected across its primary winding, which is now functioning as the secondary of a high-burden current transformer. The voltage across $R1$ is rectified by $D2$ to produce a usefully high voltage across $C1$, which will be proportional to $T4$ primary current and approximately proportional to klystron-filament current. The value of $R1$ will be reflected into the $T4$ primary circuit as a resistance of $R1/n^2$, where n is the $T3$ turns ratio. The voltage developed across the $T3$ secondary, which serves as the current-transformer primary, will be $I_p \times R1/n^2$, and the voltage across $R1$ will be

$$I_p \times \frac{R1}{n^2} \times n = I_p \times \frac{R1}{n},$$

where I_p is the $T4$ primary current. If, for instance, the klystron heater power is approximately 500 W and the $T4$ primary voltage is 120 V, the $T4$ primary current will be around 4 A. If a voltage of 5 V across $R1$ is adequate to operate the sensing circuit and $T3$ has a 120-V nominal primary and a 5-V center-tap secondary, the $n = 120 \text{ V}/2.5 \text{ V}$, or 48. The required value of $R1$, then, is $48 \times 5 \text{ V}/4 \text{ A}$, or 60 ohms. The voltage drop introduced into the primary circuit is $5 \text{ V}/48$, or 0.1 V, and the power dissipated is about 0.4 W.

The filament circuit also limits cold-filament inrush current. This feature is provided by the series ac-rated capacitor, $C2$. The cold resistance of a filament is typically one-fifth that of a fully heated filament. The starting inrush current, however, is usually specified as no more than twice the normal operating current.

Additional series impedance is required for start-up. For this purpose, designers often use a series resistor that is later short-circuited by relay contacts. An alternative that consumes no power is to use a series capacitor. A value can usually be determined that will limit inrush current to 1.5 times normal current. At turn-on, its capacitive reactance will be the current-limiting impedance. However, when the heater warms up, its resistance will increase by a factor of 5 and it will be the predominant impedance. The capacitive reactance, which is in quadrature with the resistance, will then have only a minor influence on total impedance. However, it will produce a small loss in total primary voltage, which is compensated for by the adjustable boost voltage provided by variable autotransformer *T1* and step-down transformer *T2*, whose secondary is connected, series-aiding, in the primary circuit. Meters *M6* and *M7* can be calibrated to produce quite accurate indications of actual klystron-filament current and voltage, especially near the nominal operating conditions.

Comparator circuit *U4* is the same as *U2* and *U3*. Its input is sensitive only to current in the klystron-filament circuit, which produces ambiguous information by itself because it cannot distinguish between a normal filament and a short-circuited one. In some cases, a separate circuit that is sensitive to *T4* primary voltage might be necessary. Open filaments, however, are a far more likely occurrence than short-circuited ones, making voltage sensing needlessly redundant in most cases.

Once *U4* has sensed proper filament current, its output goes low, pulling the negative input of comparator *U7* low with it and turning off its output transistor *Q4*. This allows current through *R2* and *R3* to charge *C3*. The time-constant $(R2 + R3) \times C3$ is made approximately equal to the desired cathode-heating time, which can be as great as 15 minutes, or 9000 seconds. The voltage across *C3* is applied to the negative input of comparator *U8*. The adjustable voltage at the positive input of *U8* is set to equal the *C3* voltage at the end of the desired heating interval, at which instant the output of *U8* goes low, again enabling current through the optoisolator photodiode and the green LED indicator. This action completes the path through the series-connected optoisolator phototransistors up to this point.

Resistor *R3* in the timer circuit provides some memory in case of a momentary interruption of filament power that would cause the *U7* output to conduct to ground again. If *R3* is the same value as *R2*, *C3* will discharge at an initial rate that is only twice as great as the one at which it was charged. If the filament-power interruption is only a small portion of the total heating time, the cathode will not have cooled off that much and the timing circuit will not have completely discharged either.

The completion of this simplified control ladder is shown in Fig. 20-4. The last rung of the ladder represents the personnel-safety interlocks. As will be seen, the response to these very-important sensors is redundant in different ways. The first level of redundancy is not always permissible under some system-engineering guidelines. As shown, this method requires that the 120-Vac actuating power for the three-phase vacuum contactor, which makes and breaks the primary-power feed to the high-voltage dc power supply, pass through the contacts of the personnel-safety interlock switches. These interlocks monitor both

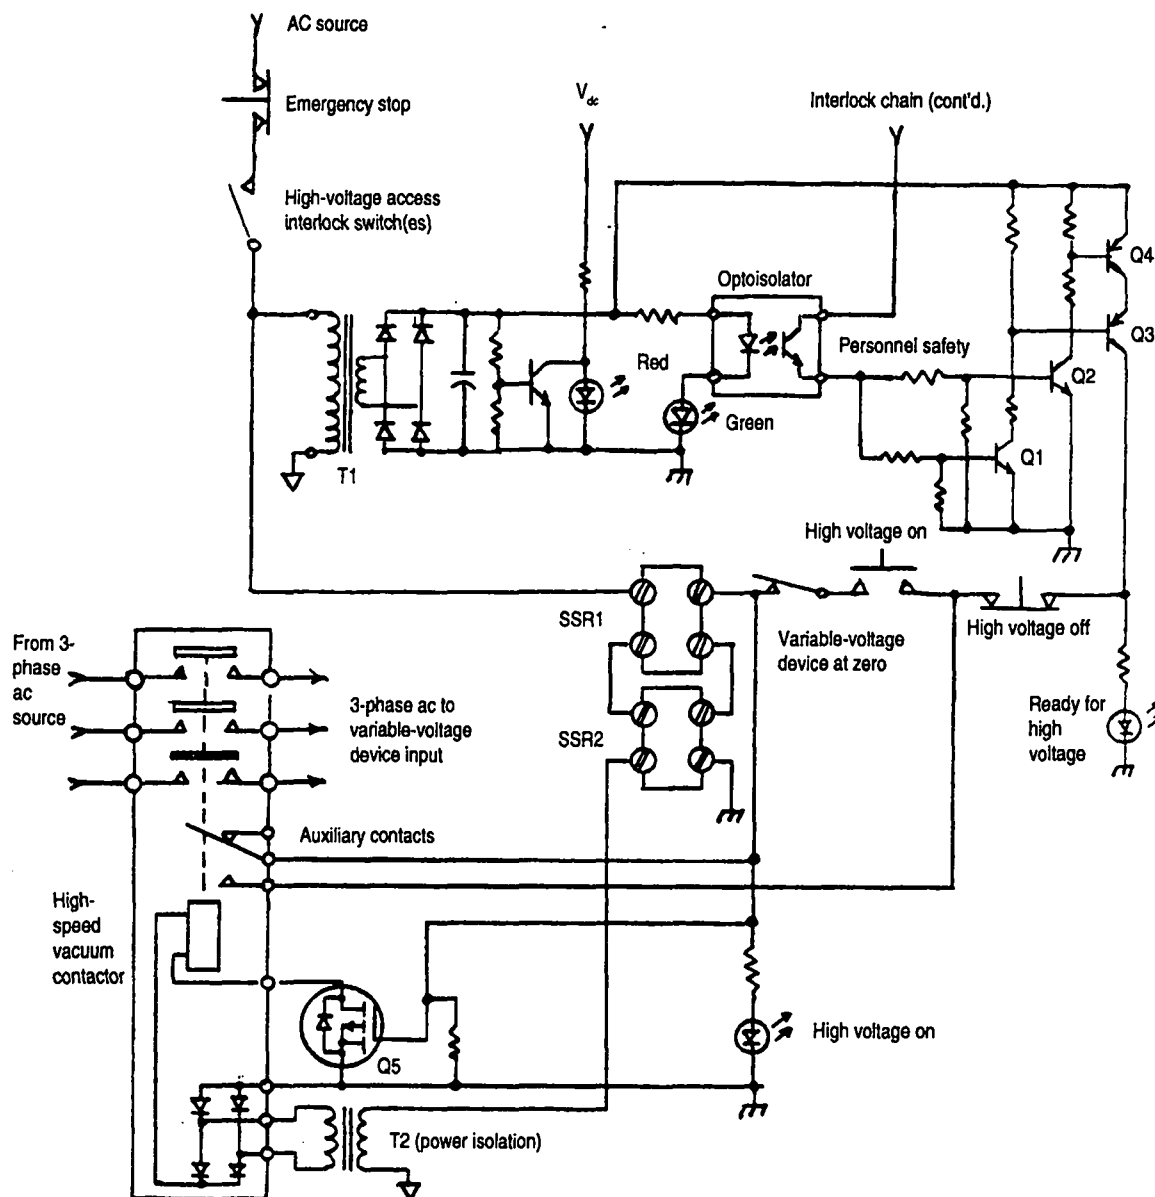


Figure 20-4. Completion of solid-state, serial/parallel control.

the status of panels and doors that permit access to hazardous voltages and the emergency-stop push button(s). This configuration is the most direct and positive means of assuring shutdown. When used, it is often augmented by running the same line to an undervoltage-trip coil on the three-phase-source main circuit breaker, making the primary-power break doubly redundant.

The objection to this approach is that the safety interlocks carry 120 Vac and therefore are potentially hazardous themselves. In addition, their closed contacts can carry significant load current, depending upon the closing and hold-in volt-ampere demands of the vacuum-contactor. Therefore, they are subject to arcing and possible contact welding. (The first objection is usually governed by some regulation, and not much can be done about it if 120-V interlocks are proscribed. The second objection can be surmounted by design. If cabinet inter-

locks like those illustrated in Fig. 4-2 are used, contact welding is not much of a threat, and positive action can easily be verified.)

If the use of 120-V interlock power is permitted and deemed advantageous, it can also be used to provide the last level of dc control power as well, after being stepped-down by $T1$ and rectified and filtered. Before the interlocks are complete, or "true," control-power V_{dc} provides the current through the red LED that indicates the incomplete condition. When the interlocks are completed, rectified dc turns on the transistor shunting the red LED, diverting current from it. The transistor also produces current flow through the optoisolator photodiode and the green LED, which indicates that the interlocks are complete. When the last optoisolator phototransistor conducts, it completes the series-connected AND gate. This connection enables current flow into the bases of transistors $Q1$ and $Q2$, which are part of a self-redundant output switch. Turning $Q1$ on turns on $Q4$, and turning $Q2$ on turns on $Q3$. Transistors $Q3$ and $Q4$ must both be on before the locally derived control dc reaches the high-voltage on/off circuit. This state is signaled when the green "ready for high-voltage" LED is illuminated. No single-point failure will cause the output switch to malfunction in an unsafe way.

If 120-Vac interlocks are not permitted for whatever reason, V_{dc} replaces the ac source, and step-down transformer $T1$ and its output rectifier are not needed. The ac is connected directly to the ac line-side of $SSR1$. With all interlocks complete, high-voltage may be turned on by depressing the high-voltage-on push button, but only as long as the variable-voltage device is at zero (as before) and its lower-limit switch is closed, as shown. Also for the sake of redundancy, contactor-actuating ac is connected to power-isolation transformer $T2$ through two series-connected solid-state relays, $SSR1$ and $SSR2$. The SSR dc control voltage is also applied to the gate of a high-voltage MOSFET, $Q5$, which is in series with the return lead of the internal rectifier for the vacuum-contactor actuator. This not only effects redundancy, but it greatly speeds the opening time of the contactor by immediately interrupting the solenoid current. In so doing the MOSFET must withstand a severe voltage transient that will appear across the dc solenoid as its internal inductance attempts to sustain current flow to keep the contactor closed.

Like the daisy-chain interlock ladder, interruption of any of the monitored and interlocked functions of the circuit just described will cause the same direct interruption of the high-voltage power supply, except that now if only one function is to blame, only one green LED will be off and only one red LED will be on. If others functions share the blame, they will reveal themselves in the same way.

20.3 Integrated-circuit considerations

The designers of modern, hard-wired transmitter-control circuits have been able to take advantage of an ever-increasing repertoire of computer-logic integrated circuits to accomplish increasingly ambitious performance goals. The only such device used in the simple control circuit just described is a dual-input NAND gate. Two of them ($U5$ and $U6$ in Fig. 20-3) were interconnected to create a set/reset flip-flop. A negative-going voltage change at the S gate input to $U5$ will shut down the transmitter. It is the intent of the design to limit such a voltage change only to excessive peak body current in the klystron RF amplifier

tube. However, what about the "noise immunity" of that input? Are there other transient voltages coupled to that point that, one way or another, could also trip the circuit and shut down the transmitter? There are events that are completely systematic, such as the closing of the three-phase primary-power contactor, that are capable of generating large-amplitude transient voltages and currents. What assurance do we have that such a transient will not trip our flip-flop or any other similar circuit element that we have designed into our system? Unfortunately, the answer is that we really have no assurance until the transmitter is completed and tested.

There are, however, choices that will either increase or decrease the likelihood of false alarms. Perhaps the most popular of all logic-element families is TTL, or transistor-transistor logic. At its most basic, TTL uses two transistors. The collector of one transistor is connected to the base of the next. When the first transistor is turned on, its collector voltage is lower than the base voltage required to turn on the next transistor. The collector of the first transistor is connected to the supply rail, nominally 5 Vdc, through a pull-up resistor. When the first transistor is turned off, its collector voltage is pulled up to a voltage that exceeds the base voltage required to turn on the next transistor. Thus, transitions are made from logical 0 to logical 1.

In Fig. 20-5 these voltage relationships are shown in terms of specification limits for TTL logic. Throughout the family of devices, the collector voltage of no conducting transistor will exceed 0.4 V and the base voltage required to turn on any transistor will not be less than 0.8 Volts. There is thus a noise-voltage margin of the difference between the two, or 0.4 V. At the other end, the maximum base voltage required to turn on any transistor is 2 V and the minimum collector voltage of any non-conducting transistor is 2.4 V. The noise-voltage margin, or noise-immunity, is again the difference between the two, or 0.4 V. Many interfaces between transistors will have higher actual margins, but they cannot be guaranteed by specification.

Integrated-circuit manufacturers have not been entirely unaware of the relative "hair-trigger" thresholds inherent in TTL logic because this problem can become intolerable when their devices must be used in threatening electromagnetic environments. Therefore, they have developed families of integrated circuits called high-threshold-logic (HTL) and high-noise-immunity logic (HINIL). As shown in Fig. 20-5, an additional noise-voltage margin was introduced by inserting a series Zener diode between input and transistor base, and by operating the circuit at higher rail voltage, either 12 Vdc or 15 Vdc, as illustrated by the simple dual-input NAND gate. Operating at 15 Vdc, the noise immunity was increased to 3.2 V for the low-to-high transition, and 6.5 V for the high-to-low transition.

These families of integrated circuits are not as fast as TTL to be sure, but what's the big hurry? In the case of the transmitter control previously discussed, the fastest fault-corrective event—the opening of the three-phase contactor—will require on the order of one-half cycle at the primary-power frequency (typically 60 Hz), which is 8.3 ms, or 8,300 μ s. It has been suggested that there is no good reason to have an individual response time that is less than one-tenth that of the ultimate response, which means that the toggle time of the flip-flop could be 830

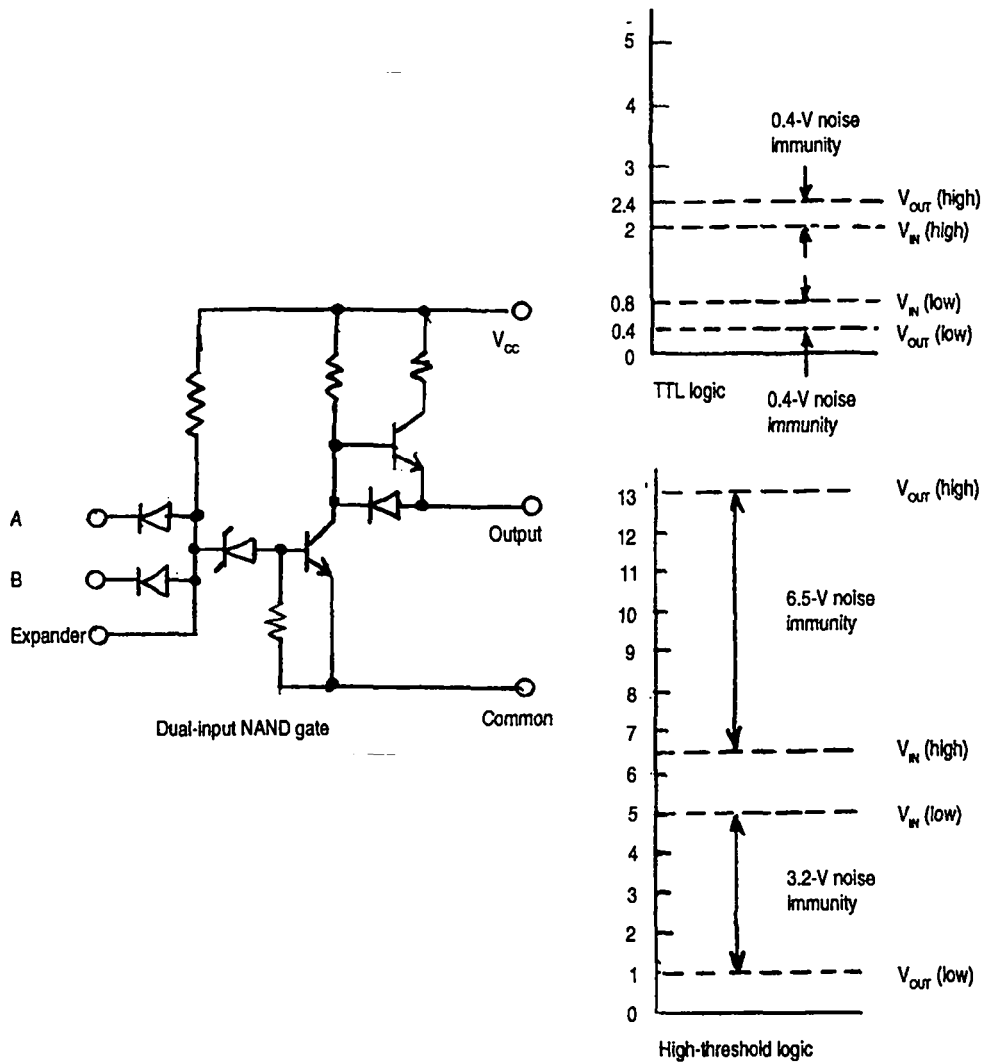


Figure 20-5. High-noise-immunity (high-threshold) integrated circuits.

μ s without noticeable effect on the overall response time. Making the response faster only increases the bandwidth of the transient spectrum to which the system might respond. With this in mind, designers have been known to deliberately slow down a NAND gate like the one shown by connecting an integrating capacitor between the "expander" input and ground. This allows the input to be present for a longer period of time before the decision is made that a real fault has occurred. (Transmitter-control logic must be slow to anger, but it should shoot to kill.)

Even where considerably faster response is required—such as when an electronic crowbar is used to discharge an energy-storage capacitor bank—submicrosecond response time is still not required. So, you might ask, why don't all hard-wired transmitter-control circuits use HTL or HINIL logic elements? Many successful ones have, but the product lines are dwindling and there are no guarantees that replacement devices will always be available. Some designers are turning to the voltage-comparator IC, from which numerous logic elements can be constructed. These devices have even greater noise-immunity levels than

HTL or HINIL, and they offer a greater promise of being around for the foreseeable future.

Regardless of what logic devices are used, their gate inputs need external protection if they are to survive for long. Figure 20-6 shows some of the protection techniques. A series resistor is always required, either to limit the fault current that a Zener-diode clamp will have to handle or to provide its share of the time-constant when an integrator or low-pass-filter capacitor is used. The diode, used as a supply-rail clamp, will prevent the gate from ever being above the supply rail by more than a diode forward-voltage drop.

A final note has to do with what could be called "connectivity safety." If we review the circuit topology in Figs. 20-3 through 20-5, we should notice that some of the functions could be accomplished with fewer active components. However, this circuit was designed so that if *any* of the removable components like the ICs or transistors were not in place, a "true" indication could not be obtained. (In the case of the *U5-U6* flip-flop in Fig. 20-3, both were part of the same IC.) Designing for connectivity safety is a good idea. At least one control circuit that did not follow the concept of connectivity safety was capable of closing the prime-power contactor whenever a circuit card was pulled from its nest. Nevertheless, connectivity safety is not the same as "fail-safe," which can never be absolutely assured.

20.4. The programmable-logic controller

The more complex a transmitter becomes and the greater the desire for its fully automatic operation, the more sense it makes to base a control system on a programmable-logic controller (PLC), its associated input and output modules, and its display and operator-input components—to say nothing of its ability to interact with other computer-based systems. Components of PLC systems have been engineered for industrial applications. They have built-in transient-voltage tolerance, hard-wired and digital filters, and with many other hardening features.

To the transmitter hardware designer, there should be nothing intimidating or

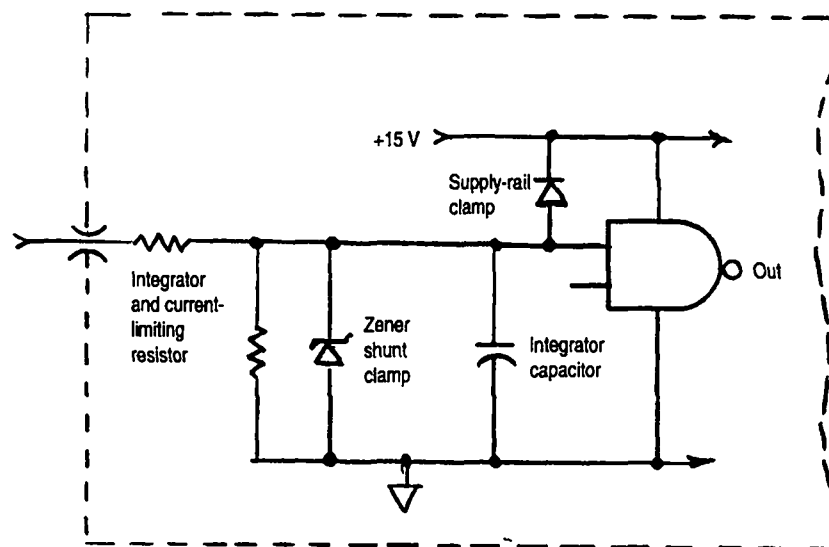


Figure 20-6. Methods of protecting integrated-circuit inputs.

off-putting about the use of a control system that will ultimately be driven by a computer program. The hard-wired interfaces are much like those that would be encountered in a conventional, dedicated transmitter-control system, such as the types already discussed. Figure 20-7 shows the layout drawing of a controller crate that is configured to handle a dual-klystron transmitter. The far-left crate location, which has no slot number, is reserved for the processor itself. A typical processor can handle over 500 input/output channels and is programmed to emulate the way that a well-designed hard-wired control circuit would respond. A major advantage of the processor, however, is that its behavior can be improved by merely changing a program. The next slot in the crate, Slot 0, is a dedicated fiber-optic converter. Two such converters can communicate with each other over either two fiber-optic channels or two twin-axial electrical cables. Slots 1 through 11 are used by different input and output modules, as shown. In this specific example, not all of the available slots are used—nor are all of the different module types used.

The dc-input modules used in Slots 1, 3, and 9 have a logic function similar to that of the relay coils in relay-coupled control logic. Each module has 16 inputs. A minimum input of 10 V is required for a low-to-high transition, a maximum of 6 V will cause a high-to-low transition, and a maximum of 30 V can be tolerated, making them even more noise-tolerant than the high-level-logic ICs previously mentioned. In addition, an input filter delays "on" response by 6 ms, and "off" response by 20 ms, and the actual coupling to the internal circuitry is by means of

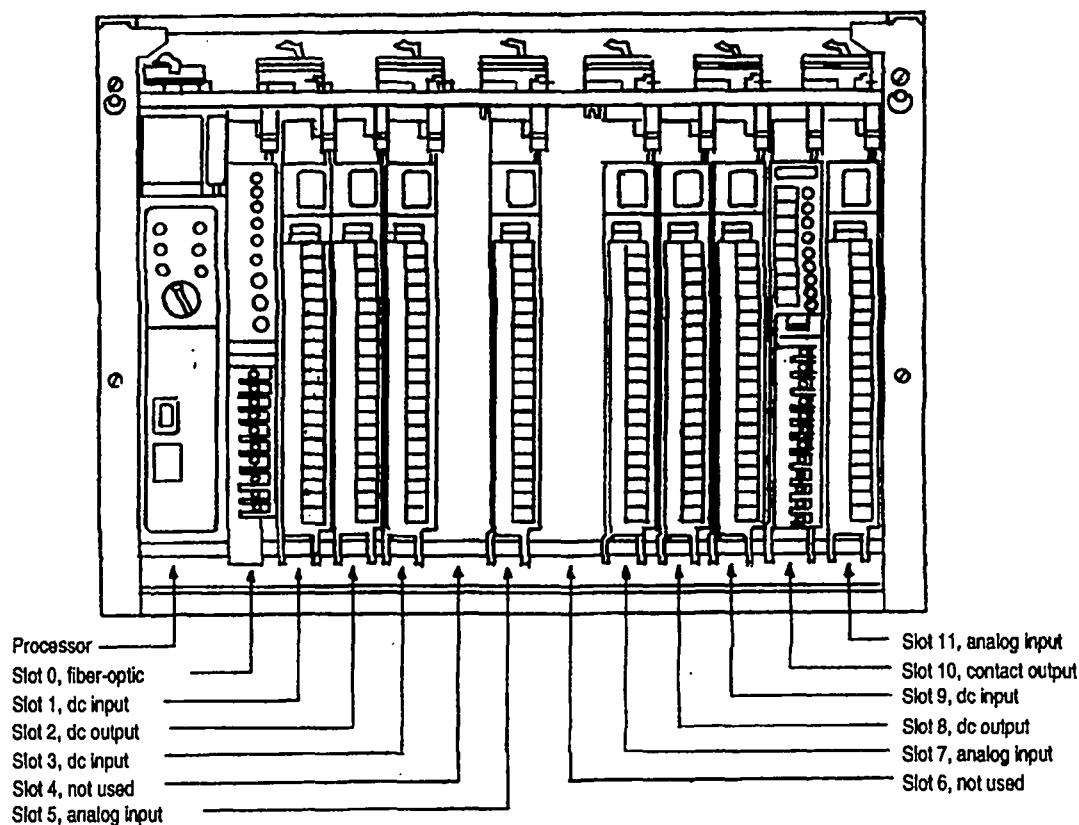


Figure 20-7. Typical arrangement of modules in programmable-logic-controller input/output chassis.

optically coupled isolators with 1.5-kV isolation.

The analog input modules used in Slots 5, 7, and 11 serve the same roles as analog or digital meters or adjustable-threshold comparators. Each will accommodate either 16 single-ended signals or eight differentially driven signals. The inputs, having a maximum range of -10 V to +10 V, are converted into either a four-digit binary-coded decimal or a 12-bit binary value, which gives the signals a resolution of 1 part in 4095. The signals are then passed on to the processor.

The dc-output modules used in Slots 2 and 8 have 16 outputs each. The input to each module is a dc power supply, which ranges from 10 V to 60 V. As commanded by the processor, each module will switch as much as 2 A per output (8 A total per module) to a load connected to an output.

The selectable-contact output module used in Slot 10 has four pairs of relay contacts, each pair sharing a common lead. Each relay contact is jumper-selectable as either normally closed or normally open. It will switch ac or dc in the range from 24 V to 125 V at up to 1 A, with 5-ms operate-and-release times. Its maximum actuation rate is 10 Hz.

Each module has its own set of diagnostic LED indicator lamps to locally indicate activity. The signal lines that the module inputs and outputs are connected to are shown in Fig. 20-8. Most of the signal designations should be self-explanatory. Some, however, may be more scrutable than others. Together they are sufficient to control the operation of a dual-klystron, pulse-modulated, high-power RF source.

Programming the processor to deal with the input and output signals in the same fashion as a hard-wired control circuit is only part of what can be accomplished. The fact that the operator/transmitter interface can be ergonomically designed (or programmed) is equally, if not more, important. Figure 20-9 shows a typical (and now almost obsolete) operator interface. It has a 12-in. color-video display, through which the control system can communicate with the operator, and 16 push buttons (or "programmable function keys"), through which the operator can communicate with the control system. Illustrated in the figure is but one of many display screens on which transmitter status is alpha-numerically presented. In addition to standard boxes—"go" (green, "OK"), "no-go" (red, "NOK")—to the right of each function, analog information is digitally presented,

DC input	DC output	DC input	Analog input	Analog input	DC output	DC input	Select. Contact	An. input
Not used	+dc	Not used	01 Interpulse mod V	10 Ky 1 fil V	+dc	Not used	1 Reset	11 Cap V
Not used	+dc	Not used	02 Ky 1 l cathode	11 Return	Not used	Not used	2 +24 V	12 Return
Not used	+dc	Not used	03 Ky 1 l body/MA	20 Ky 1 fil I	+dc	Not used	3 HV on	21 PS I
Not used	+dc	Not used	04 Ky 2 l cathode	21 Return	+dc	Not used	4 HV off	22 Return
00 Oil low	02 Focus coil on	02 Ky 1 pressure	05 Common	Common	02 HV dump up	02 CB 10 kV OK	5 +24 V	23 Common
01 Oil high	01 Fil 1 on	03 Ky 2 pressure	06 Intrapulse mod V	30 Ky 2 fil V	03 Crowbar test	03 CB 450V OK	6 HV off	24 PS V
02 Tank lid	02 Fil 2 on	04 Ky 2 l body/MA	07 Ky 2 l body/MA	22 Return	04 Crowbar on	04 HV dump down	7 Ready for HV	25 Return
03 Coil flow	03 Modulator on	05 Ky 2 l cathode	08 Ky 1 pressure	31 Ky 2 fil I	05 Return	05 HV dump up	8 HV off	26 Common
04 Coil temp	04 Rack select	06 Crowbar test	09 Ky 2 pressure	21 Return	06 Return	06 CB test rly down	9 HV off	27 PS V
05 Body 1 flow	05 Fault reset	07 Ky 1 l body/MA	10 Common	Common	07 Return	07 Gnd rly up	10 HV off	28 Common
06 Body 1 temp	06 Switch ac on	08 Ky 1 l cathode	11 Common	30 Focus coil 1 I	08 Return	08 Cap m intlock	11 HV off	29 Common
07 Body 2 flow	07 RF amp 1 on	09 Pulse timing	12 Common	31 Return	09 Return	09 Gnd stk OK	12 HV on	30 Common
08 Body 2 temp	08 RF amp 2 on	10 CB driver OK	13 Common	40 Focus coil 1 V	10 Return	10 HV off	13 Local mode	31 Common
09 Focus coil flow	09 Focus coil temp	11 Pulse deck OK	14 Common	41 Return	11 Return	11 HV on	14 Remote mode	32 Common
10 Fil 1 select	10 Fil 2 select	12 RF 1 OK	15 Common	Common	12 Return	12 Local mode	15 HVPS rem intlock	33 Common
11 RF load flow	11 Crowbar test in	13 RF 2 OK	16 Common	20 Focus coil 2 I	13 Return	13 Remote mode	16 HVPS ready	34 Common
12 Ground	12 Ground	14 Ground	17 Common	31 Focus coil 2 V	14 Return	14 CB test fail	17 Not used	35 Common
			18 Common	22 Return	15 Return	15 HVPS ready		36 Common
			19 Common	Common	16 Ground	16 Ground		37 Common
			20 Common	Common				38 Common
			21 Common	Common				39 Common

Figure 20-8. Actual inputs and outputs of programmable-controller modules for dual-klystron transmitter.

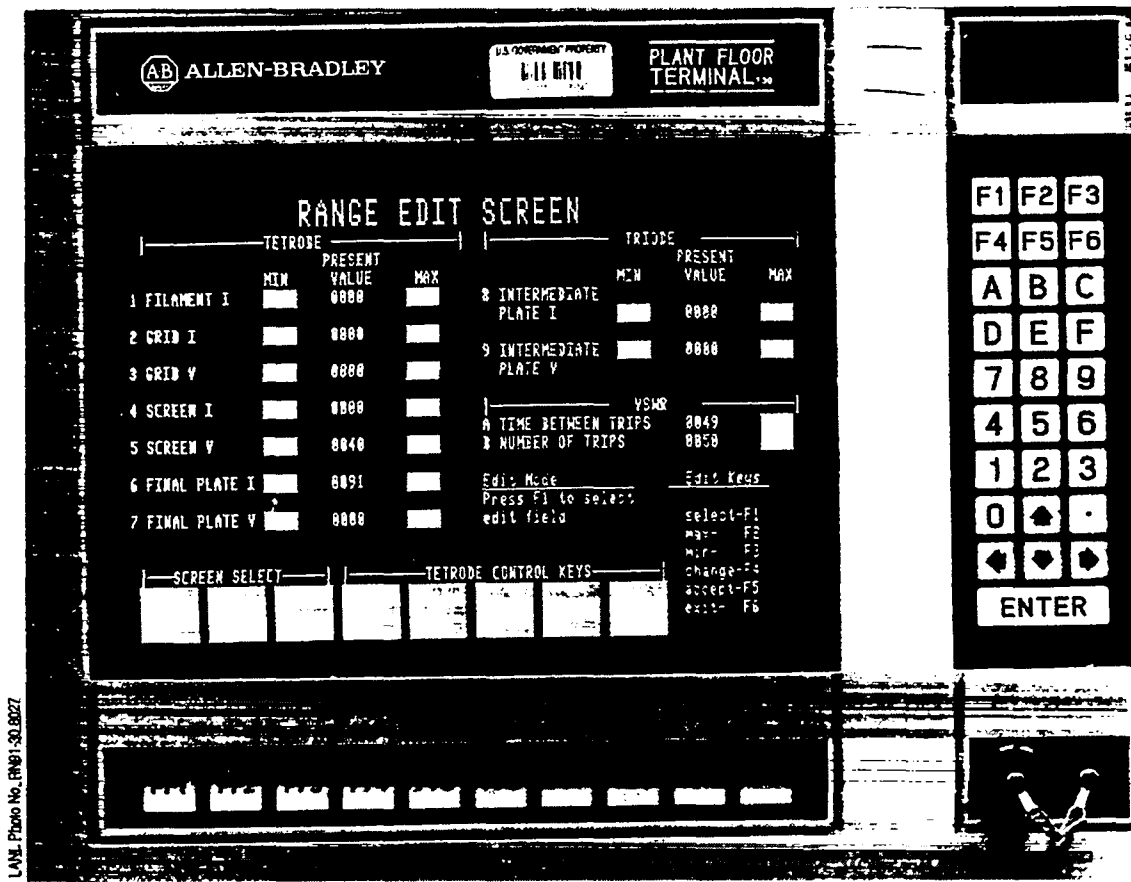


Figure 20-9. Typical operator interface for transmitter control using programmable logic.

as well as the limits to which it may be compared. The klystron-beam voltage, for instance, is presented both as a value from 0 to 100 kV and as a solid horizontal bar graph. Filament voltage and current are compared with maxima and minima, which are automatically time-compensated for the 150% cold-filament-inrush current. The cathode-heating timer can be viewed as it runs down from 900 seconds to zero. Water-flow interlock switches have digital delays associated with them, so that a momentary drop-out will not dump the entire system until approximately 5 seconds has elapsed.

More modern implementations of PLC transmitter-control systems use special-purpose software to allow even more elaborate ergonomics, including a touch-screen interface and almost cartoonlike identification of functions. But the PLC approach offers a level of dynamism that can, if not carefully controlled, get out of hand. As easy as it is to program a function and its display response, it is just as easy to change it. And almost anyone can do it if they have a computer and a keypad. Configuration control for PLCs can become a nightmare.

The system illustrated was eventually married to a host computer that was programmed to dictate its behavior. Under fully automated computer control, the PLC changed beam voltage, peak beam current, RF power level and frequency, and even the phase and amplitude of RF load mismatch, making possible the fully automatic creation of RF bandwidth plots and Reike diagrams.

To many engineers, PLC transmitter control is the ultimate, and the concept

gains proselytes easily (especially among younger technologists). The PLC will not do everything, however. In the first place it, is relatively slow. The processor does not simultaneously look at the inputs all of the time; it looks at them sequentially. It updates its outputs sequentially as well, and has a cycle time of approximately 10 ms. This is fast enough for many stimuli and responses but not for all. If the transmitter has an electronic crowbar, then it must be fired by dedicated and relatively high-speed circuits. The same is true of RF arc detectors, which should be capable of interrupting RF drive with a submicrosecond response time. Personnel-safety interlocks must also be dedicated and hard-wired in the fashion of direct interrupts. All of this additional information can be sent to the PLC, and it can respond to it redundantly, which is all the better.

20.5 High-speed instrumentation for transmitter monitoring

In order to generate transmitter-control decisions in the microsecond time domain and view on an oscilloscope pulsed currents and rapidly varying voltages that may have pulse durations and rise-and-fall intervals measured in nanoseconds, special components and assemblies are required. The two most important are the wide-band current-monitor transformer and the frequency-compensated voltage divider.

Figure 20-10 shows a simplified high-voltage circuit for a transmitter module using eight parallel-connected TWTs, each with dual-depressed collectors. To accurately view the pulse currents in the cathode and collector leads of an individual TWT requires the use of current-monitor transformers, which not only have a high degree of waveform fidelity but are compatible with the high volt-

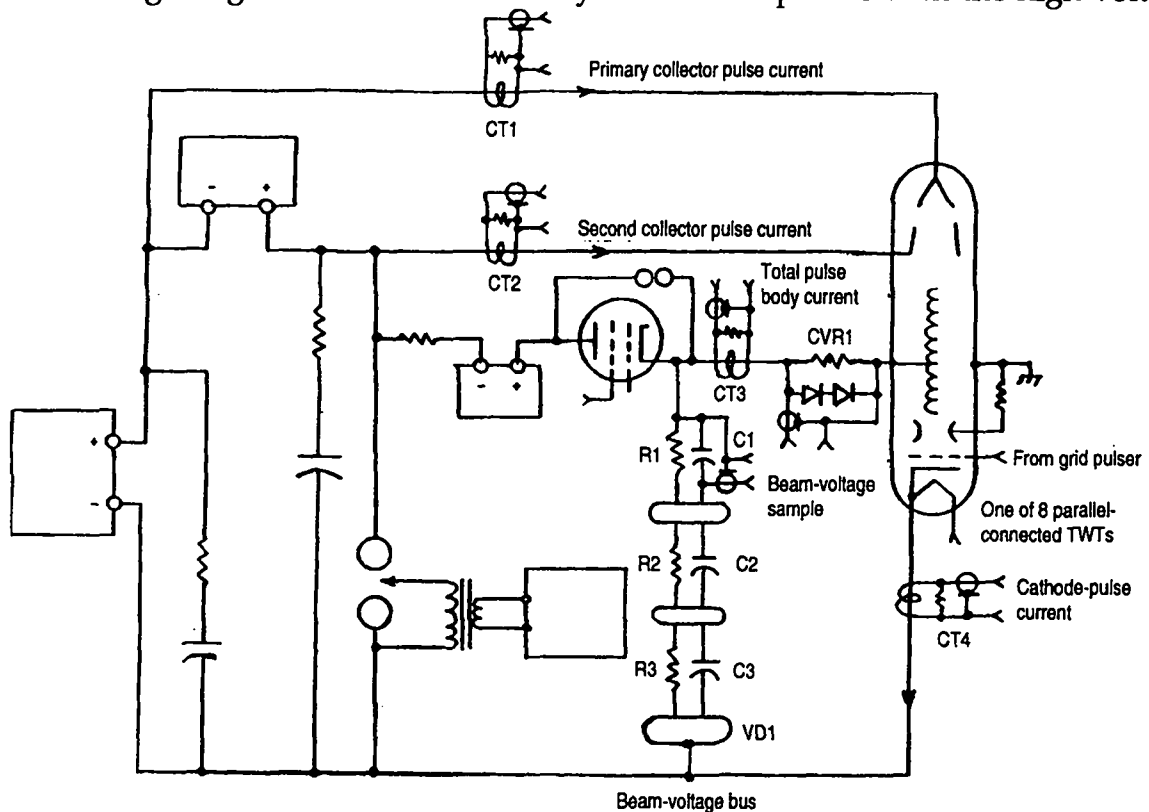


Figure 20-10. Current and voltage instrumentation for dual-collector TWT in multiple-tube transmitter.

ages applied to these electrodes. The ring-type current-monitor-transformer family—the members of which either look like doughnuts, flat-sided doughnuts, or a square metal box of modest depth with a large hole through it—is ideally suited to perform this task.

Inside the doughnut shape is an electrostatically shielded toroidal winding, one end of which is connected to the outer shield of a coaxial cable and the other end to a terminating resistor and the center conductor of a coaxial output connector. The shield, which is the outside of the doughnut, is actually in two parts. This is because it cannot be continuous, or else it would function as a short-circuited turn. The two half-shields overlap but do not touch. The conductor carrying the pulse current to be monitored is passed through the hole in the doughnut. A time-varying magnetic field surrounds the conductor when pulsed current passes through it. The rise time of the output from the toroidal winding is shorter than the delay time of the helical transmission line that the winding forms with the surrounding shield. This is due to the fact that voltage is induced into all of the toroidal turns simultaneously because they surround the current-carrying wire. The highest-performance versions of these transformers have distributed loading around the periphery of the winding to maximize the bandwidth. The low-frequency, or pulse-droop, performance is determined by the ringlike magnetic core, around which the toroid is wound. Such transformers not only have pulse droop but a volt-time product (or more correctly they have a primary current-time product limit beyond which they saturate and suddenly force the output to drop to zero). The most broad band of the transformers, those with the lowest product of useful pulse rise-time and flat-top droop rate, tend to have the smallest hole diameters and the lowest sensitivity. Useful rise times are as low as 10 ns and droop rates as low as 0.5% per ms. Typical voltage/current sensitivities range from 0.01 V/A up to 1 V/A, although very large transformers have been built to handle pulse currents in the 100,000-A range. Although these huge transformers have droop rates of 0.1%/ms, their sensitivity is only 0.1 mV/A.

The size of the hole in the doughnut is usually dictated by the voltage on the conductor, or "primary," that passes through it, and there is a wide range of hole sizes, up to a standard maximum of 10-3/4 in. Because of the coaxial geometry between wire and hole, the maximum voltage gradient will be at the outer surface of the wire. There is an optimum diameter for this conductor to minimize that gradient. It is the hole diameter divided by e (the base of natural logarithms). For the 10-3/4-in. hole, the optimum conductor outer diameter is 4 in. and the flashover voltage for sea-level air is 150 kV.

Most current-monitor transformers have 50-ohm internal impedances. When feeding a standard 1-Mohm oscilloscope-input resistance through a cable of 50-ohm characteristic impedance, the voltage sensitivity is as rated. The scope impedance produces a 100% voltage-aiding reflection, which is totally absorbed by the transformer internal impedance without re-reflection. When the scope end of the cable is terminated in a matching resistance, the sensitivity is cut in half. There is no practical advantage to terminating the scope or load end. Not all transformers have 50-ohm internal impedance. It is possible to combine higher sensitivity and lower droop rate if the internal impedance can be made greater.

For long-duration pulses of modest rise time, the mismatch will produce pulse distortion that is usually not noticeable. Values up to 500 ohms have been used successfully. The transformer output voltage, which is produced across an actual internal impedance, represents finite power that must be supplied by the primary circuit. This means that there is a finite resistance coupled into the primary conductor. It is usually of negligible magnitude, but it can be as great as 0.2 ohms.

Although the output of a high-quality current-monitor transformer can be a highly accurate replica of the sampled current, it can never have all of the properties of that current. In the first place, the monitor is a transformer and, therefore, can have no average value of voltage across its output. In a rectangular-pulse system, it is restricted to duty factors less than 50%. Furthermore, even in low-duty pulse systems, the baseline can never be zero volts. The area of the volt-time product above zero must be the same as that below zero. In a 10%-duty-factor system, for instance, the pulse itself may be reproduced with great fidelity, but if the transformer droop rate is small compared with the interpulse interval, it will start at a voltage that is negative with respect to ground by 10% of the pulse amplitude and return to that level when the pulse ends. This offset will increase as duty factor increases until it reaches a maximum value of 50% at 50%-duty factor. If the transformer droop rate is small compared with the interpulse interval, each output pulse will start at zero, the pulse top will droop (even if imperceptibly), and the pulse will under-shoot the baseline at the end of the pulse, exponentially recovering to zero before the next pulse. The requirement for equal volt-time areas will be met, but with an exponentially decaying negative voltage.

Returning to Fig. 20-10, note that the "body," or ground, current is monitored by both a current-viewing resistor, CVR1, and a current-monitor transformer, CT3. However, in this example there are eight tubes in parallel whose bodies must all be grounded. The current monitored, therefore, is the total body current of all of them.

The broad-band voltage monitor shown is a frequency-compensated voltage divider, VD1. In this example, it is used to view the TWT beam-voltage bus, nominally -40 kVdc. The output sample across R1 and C1 is typically in the range between 1/10,000 and 1/1000 of the bus voltage. The compensated divider is not only a resistive divider comprising R1, R2, and R3 (and having a frequency response down to zero frequency) but a capacitive divider comprising C1, C2, and C3 (and having theoretical frequency response of everything but zero frequency). If the two are connected in parallel and have the same division ratio, their combination has a theoretical total frequency response from dc to infinity. The equal division ratio is the same as the equal R-C product: $R1 \times C1 = R2 \times C2 = R3 \times C3$. The upper-frequency limit is determined by series inductance in both the resistive and capacitive legs. Clever component design can extend frequency response to the megahertz region. Some designs even use printed parallel-plate capacitance with the resistance mounted to the printed wiring in order to minimize inductance. These designs are more successful when insulated with dielectric oil because the printed-capacitance edge effects tend to produce very high electric-field enhancement. When these devices are air-insulated, it is customary

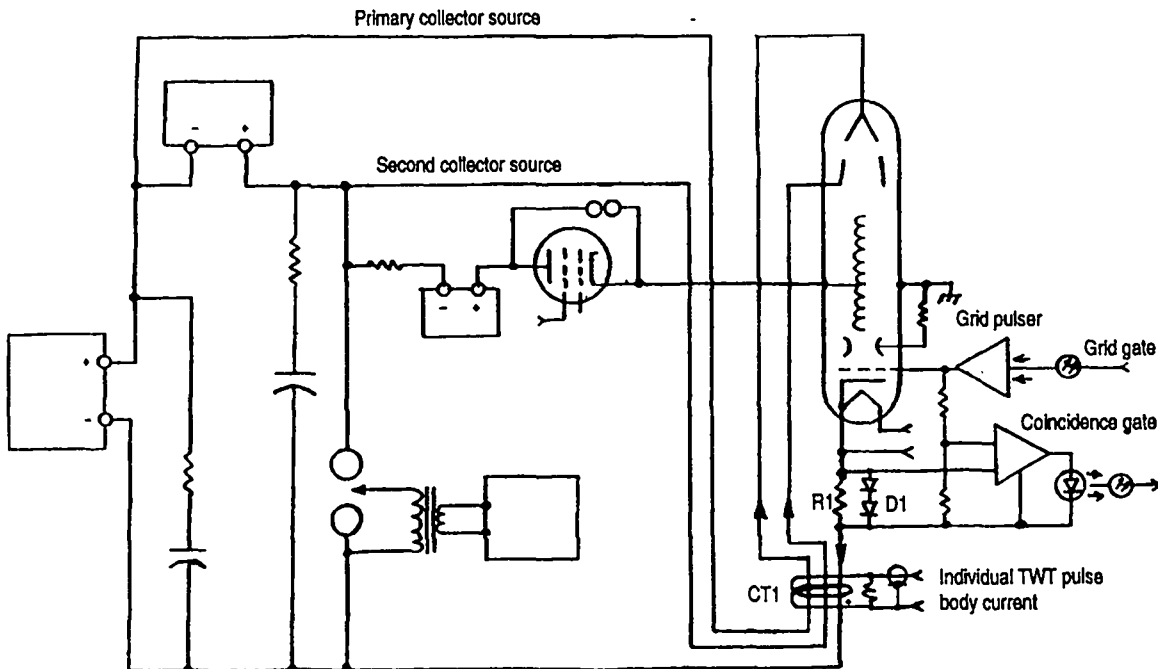


Figure 20-11. TWT of multiple-tube transmitter instrumented for individual-tube body current and interpulse cathode current.

to break up the divider into a number of layers with some corona-reducing metal shape, such as the doughnut, separating them. Two layers are schematically represented in the illustration.

A solution to the problem of monitoring body current for each individual TWT in the eight-tube transmitter is shown in Fig. 20-11. Body current in a microwave tube is current that does not go somewhere else. The "somewhere else" in the TWT illustrated is the two collectors. If the total collector current is subtracted from the cathode current, the remainder is body current (or body-plus-grid current). This subtraction can be accomplished vectorially by passing the cathode lead of each TWT through a current-monitor transformer window in one direction and the two collector leads through the same window in the opposite direction, as shown. The net magnetic field in the window will be proportional to the vector difference between the cathode and the two collector currents, which is the body current. (This is the same way that many household ground-fault-interrupter circuits work. Ground current is inferred from the difference between high-side supply current and low-side, or neutral, return current as subtracted by a current transformer.) However, the ground current indicated will also include the charging and discharging currents of the stray capacitances associated with the grid or any other modulating electrode supplied by the pulse modulator.

Although this technique of inferring body current is effective in most situations, sometimes reality is not what it appears to be. Consider a medium-power gridded TWT. It operates at very short pulse duration, typically 100 ns, and has rise-and-fall times of less than 10 ns. The beam voltage is typically 10 kV or so. The electron beam will be travelling at approximately 20% the speed of light. If the beam traverses one foot of interaction region between cathode and collector,

it will take about 5 ns for the electron beam to make the trip. The collector-current waveform will be delayed from the cathode current waveform by 5 ns, and the subtraction process will give very large errors during the rise-and-fall intervals of the pulse because two large numbers are being subtracted to get a small one. Fortunately, in most cases it is the mid-pulse body current that is the most important, and its replica will not be distorted.

Not only does vector subtraction work in a current transformer but vector addition does too. The sensitivity of a transformer can be doubled by passing the "primary" lead through the transformer window twice in the same direction, or it can be multiplied by n by passing the lead through n times. However, the resistance coupled into the primary circuit will increase by the same factor.

Sometimes in a high-power gridded microwave tube it is crucial to be able to measure even small values of cathode current, down to zero frequency, and compare them with the timing of the grid pulse. In Fig. 20-11, that is the function of $R1$. Diode $D1$ shunts the major portion of the intended pulsed beam current. The voltage across $R1$ is compared with a portion of the grid-drive pulse. If it overlaps by too much, the system is shut down.

A major threat to the longevity of such a tube is the open grid, a situation in which everything is on while the grid is floating, not connected to anything. This is not just a hypothetical condition, especially in a transmitter that has 128 such tubes. What happens if the grid is unconnected is that it will be self-biased to a point near beam-current cutoff, called grid-leak biasing. Electrons striking the grid will produce a negative potential on it. The current will be infinitesimal because the circuit impedance is infinite. Beam current will be very small but continuous—and worst of all, super-focused. It will pass down the center of the beam tunnel, much like a laser beam, and will disperse only slightly when it leaves the beam tunnel, or interaction region. This beam current is capable of boring a hole in the center of the collector, a fatal occurrence for the tube.

20.6. High-speed crowbar-firing circuits

There are nearly as many low-level crowbar-firing circuit designs as there are circuit designers. Nevertheless, the circuit shown in Fig. 20-12 illustrates some of the features that such a circuit should have. The inputs to the circuit are outputs from current-monitor transformers. Three monitor points are shown for a single microwave tube, in this case one with a modulating anode. The pulse currents sampled are cathode, body, and modulating anode. A typical 1-2 MW peak-power modulating-anode klystron will operate at a beam voltage of less than 100 kVdc with peak-pulse cathode current of 30 A, body current of 0.3 A (or 1%), and modulating-anode current of 0.03 A (0.1%). The customary transformer sensitivities used are 0.1 V/A for cathode current (3 V peak, nominal), and 1 V/A for body and modulating-anode currents (0.3 and 0.03 V peak, nominal).

A 15-V, high-surge-current Zener diode directly shunts the output of each current-monitor transformer with no intentional series resistance at the input to each channel. Its purpose is to clamp the output voltage of each monitor transformer to 15 V peak. Peak instantaneous fault current due to an electron-gun arc in a transmitter like this is typically between 1000 and 2000 A (100 kV limited by 50 to 100 ohms). It will flow through both the cathode- and body-current moni-

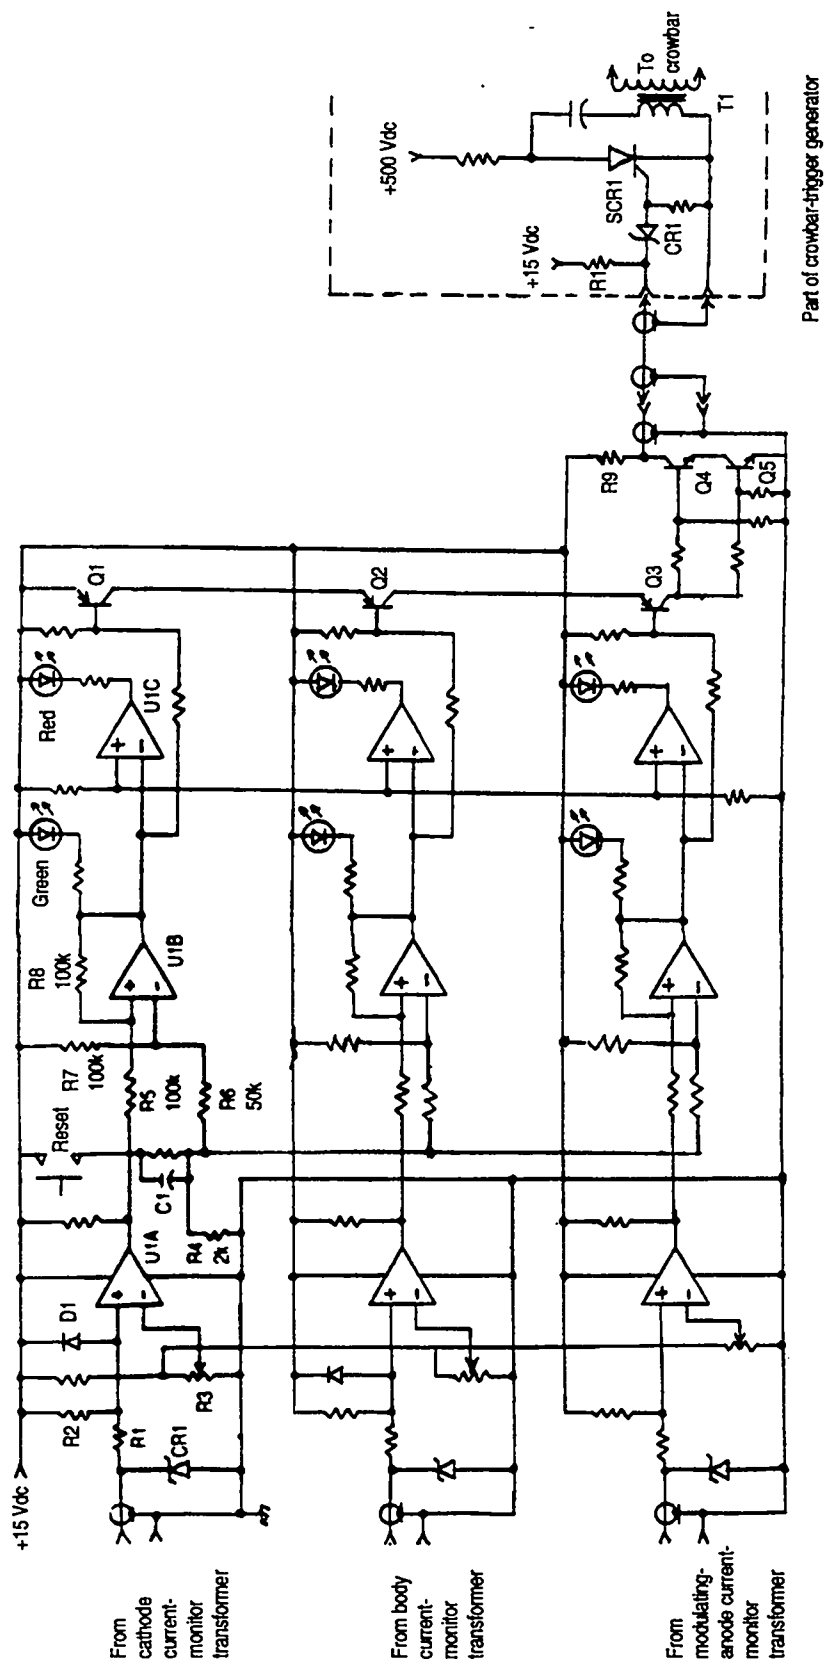


Figure 20-12. Crowbar-firing logic circuit using voltage-comparator integrated circuits.

tors. With no clamping, the peak output voltage of the cathode monitor would be between 100 and 200 V and the body-current monitor would be between 1000 and 2000 V. Neither output pulse would last very long, because both transformers would almost instantaneously saturate due to excessive $V \times t$ product. The desired input to the fault channel is no more than 15 V for a time duration that is at least long enough to insure that the circuit reacts. (The maximum peak current through the cathode-current-channel Zener diode is 200 V/50 ohms, or 4 A. The maximum peak through the body-current-channel Zener diode is 2000 V/500 ohms, or 4 A because the transformer is a non-standard high-impedance type.)

Each of the three channels is identical. The first stage is a voltage comparator, typified by *U1A*, that has an adjustable-voltage threshold set by potentiometer *R3*. Resistor *R1* limits current in the signal channel. If, for any reason, input voltage exceeds the dc rail voltage, the excess will forward-bias diode *D1* (one of the protective diodes first shown in Fig. 20-6). Resistor *R2* normally does nothing. It has a high value (100 kohm or more) compared with the very low internal impedance of the current-monitor transformer that is normally connected to each input. Should a transformer be disconnected, however, or never connected in the first place, *R2* will pull up the gate input to the rail. This will have the same effect as a massive fault. This strategy is consistent with the connectivity-safe philosophy.

Comparator *U1B* is connected as a set/reset flip-flop by means of resistors *R5* through *R8*. This use of a comparator is another high-threshold alternative to TTL logic. In the "reset," or normal, state of the flip-flop, the voltage at the negative input of *U1B* is the voltage across *R6*, which is about +5V. The voltage at the left end of *R5* is a small fraction of a volt because *U1A* is normally turned on. The output of *U1B* in the reset state is also conducting, as we will see. The left end of *R5* is low and the right end of *R8* is low also, so the voltage at their junction, which is the positive input of *U1B*, must also be low. This is a stable state.

Any input voltage exceeding the threshold voltage of *U1A* will cause its output to go from low to high. When this happens, the left end of *R5* will be pulled high (the output pull-up resistor is small compared to 100 kohm). The right end of *R8* is still low at this moment, but the voltage at the resistor junction will rise to about 7.5 V. When it passes +5 V (the voltage at the negative input of *U1B*), the output will toggle from low to high, raising the voltage at the positive input momentarily to +15 V. As soon as the input fault voltage recovers to zero, the output of *U1A* will go low again, and the voltage at the positive input of *U1B* will stabilize at +7.5 V with its output high. This is also a stable state. Manual reset is accomplished by pushing the reset switch, which momentarily pulls up the voltage at the left end of *R6* through *C1*. The voltage at the negative input of *U1B* is pulled to +15 V, toggling its output from high to low, which is where the process began.

When the output of *U1B* is low, it enables current to flow in the base of *Q1*, part of an AND gate that includes *Q2* and *Q3*. When all three are conducting, series-redundant transistors *Q4* and *Q5* conduct as well, providing a low-voltage-drop pull-down, or sink, for current in the output circuit. The output of the crowbar-firing circuit is connected to *SCR1*, the first capacitor-discharge stage of




Figure 20-13. A typical source of ground-loop, or "ground-bounce," spurious signal.

the crowbar-trigger generator. The output of the capacitor-discharge stage initiates breakdown of the high-voltage electronic crowbar. A local power supply in the trigger generator will cause gate current to SCR1 to flow through R1 and CR1, unless their junction point is pulled down by conduction of Q4 and Q5. The cable between the crowbar-firing logic and the trigger-generator must be connected, or the crowbar will fire (or the discharge capacitor will fail to charge to +500 V, a situation which is interlocked). This is more connectivity safety.

The third voltage comparator, U1C, is used as inverter to illuminate the red LED that indicates a fault in that channel. The same information is conveyed to a PLC input module or comparable interface of a generic control system.


20.7 The pulse "balun"

The last component to be discussed is the pulse "balun." The generic balun is a circuit element that allows a differential, or balanced, circuit to be connected to single-ended, or unbalanced circuit, or vice-versa. (The "bal" in its name is for balanced, and the "un" is for unbalanced.) The pulse balun is a circuit element that permits this to happen over the frequency domain associated with impulses. Its most common use is not to couple signals but to permit unwanted signals to be suppressed.


A classic example of a situation that will generate an unwanted common-mode signal is shown in Fig. 20-13. The offending circuit is part of a line-type modulator. The designer's mistake, whether intentional or not, is to use ground as the return conductor for the low-side of the high-current discharge circuit. By doing this, the designer subverts the ground; it is no longer ground. Ground, no matter how high its conductivity is intended to be—a copper-sheet floor notwithstanding—has resistance and, more importantly, inductance. Current will produce an $I \times R$ voltage drop, and change-of-current will produce an $L \times di/dt$ voltage drop between points A and B. If the discharge current is sampled by

means of a CVR in the cathode of the discharge switch, referenced to point A, and the oscilloscope used to monitor the CVR voltage is connected to safety ground, point C, there will be shield current in the interconnecting cable that will superimpose all or part of the "ground-bounce" voltage on the composite signal that the scope thinks it is looking at. (An acid test for this situation is to touch just the outer conductor of the cable to the outer surface of the scope input jack. If there is a vertical scope trace, it is time for a pulse balun.)


A pulse balun can take many shapes, but it is usually nothing more than a length of coaxial signal cable wound around a magnetic core. Figure 20-14 shows two common types. In Fig. 20-14a, the cable has been wound toroidally around a ring-like core. In Fig. 20-14b, a longer length of cable has been wound around a standard transformer iron core for greater low-frequency performance. The equiva-



A. Toroidally wound pulse balun



B. Pulse balun wound on conventional transformer core



C. Equivalent circuit of pulse balun

Figure 20-14. Aspects of the pulse balun.

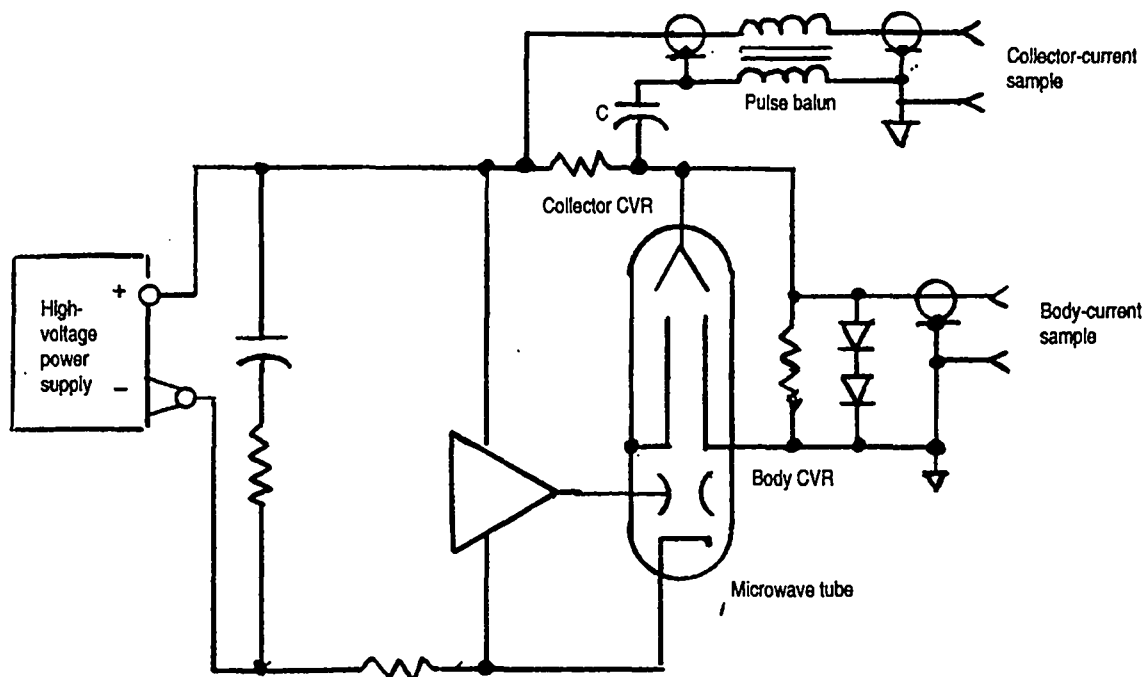


Figure 20-15. Pulse balun used to translate differential collector-current sample to a single-sided signal referenced to ground.

lent circuit is shown in Fig. 20-14c. (The pulse balun is actually the same thing as a bifilar transformer winding, but without the other winding or a bifilar inductor.) The two conductors of the cable—the inner conductor and shield—have unity coupling. Whatever voltage was between them at one end is there at the other end. The shield, however, now has series impedance. A varying voltage can exist between the shield at one end and the shield at the other end with current limited by the winding inductance.

If the cable of Fig. 20-13 had a pulse balun in series with it, the ground loop would be open-circuited by it. Shield current resulting from the differential ground voltage between A and B would be limited by the winding inductance of the balun, and the scope would respond only to the voltage between cable shield and inner conductor. There was a time in virtually all high-power installations, especially where current levels were in the multiple kiloampere range, when all coaxial-cable interconnections were made via pulse baluns.

Figure 20-15 shows another practical application of the pulse balun. It is in an application where a direct-coupled sample of collector current is required to guard against "collector boring" should the modulator fail to completely cut off beam current in the interpulse interval. Both ends of a CVR must float, however. One end is connected to the body-current CVR, and the other end is the insulated low-voltage return of the high-voltage dc power supply. The only ground reference is at the body of the microwave tube, as usual. A pulse balun can be used to develop outer-conductor winding inductance that limits the current flow that would shunt the body-current CVR. The body-current pulse-top droop rate would be determined by L/R , where L is the pulse-balun winding inductance and R is the body-current CVR value. An important detail is the coupling and dc-block-

Index

Symbols

μ factor 166
24-pulse-rectifier 272
3-dB hybrid 231
4CPW1000KB tetrode 173
4PR250C/8248 tetrode 218
7560 triode 169
8960 tetrode 218

A

Accordion resistor 262
Acoustic-shock effect 39
Adiabatic 255
ALCOR/Millimeter-Wave Radar 15
ALTAIR 32, 249
AM. *See* amplitude modulation
Amplification factors 147
Amplitron 191
Amplitude imbalance 269
Amplitude modulation 109, 111
Amplitude-pushing factor 115
Analog input modules 367
AND gate 352
Angle of retardation 283
Anode delay-time drift 126
Anode delay-time jitter 125
Anode dissipation 167
Anode firing 124
Anode-grid hold-off voltage 163
Antenna duplexer 25
Arc
 cross-guide 14
 detector
 waveguide 20
 series 14
Arrhenius equation 30
 definition 10
Attenuator
 drive-power-level 19
Availability 31

B

Back-lit thyratron 128

Back-swing clipper circuit 84, 154
Ballistic Missile Early Warning System 220
"Bandwidth," voltage 293
BANSHEE 94
Bifilar winding 104
Blumlein discharge circuit 93
BMEWS. *See* Ballistic Missile Early Warning System
"Bolted" fault
 definition 9
"Boost" 292
Breakdown voltage 55
BST (beam-switch tube) 198
"Buck" 292
Buck regulator 311
Buck/boost converter 314
Built-in test equipment 39

C

Calefaction 350
Capacitive input filter 274
Capacitive storage 152
Capacitor bank
 step start or soft start 7
Cathode-controlled electron gun 65
Cathode-heating time 360
CFA. *See* crossed-field amplifier
Characteristic impedance 73, 78
Charge diversion 152
 definition 6
Charge diverter *See* crowbar
Charge diverters 253
"Charge pump" 314
Charge-transfer 135
Charging circuit 85
Charging inductor 86
Child's law 2
Circuit breaker 308
 definition 9
Circulator 18
 high-power 22
Clamp winding 315
Clamping diode 117
Clipper diode 84
Closed-loop water-cooling systems 342

Coaxial transmission line 12
Cobra Dane radar 32
Cockroft-Walton generator 325
Cold-filament inrush current 359
Combiner
 hybrid 22
Common-mode signal 376
Commutation
 phase 266
Commutation angle 266
Commutation factor 267
Condenser 350
Conduction 334
Conduction angle 286
Configuration control 41
Control and monitor 10, 352
Control-grid modulators 234
Convection 334
Cooling system 10
Core saturation 104, 238
Corona 57
Corona discharge 159
Corona shields 159
Coupler
 forward and reverse 22
 signal 17
 waster power 23
Crossatron 187
Crossed-field amplifier 114
Current spreading time 133
Current-monitor transformer 37
Current-viewing resistor 371
CVR. *See* current-viewing resistor

D

Daisy chain 352
Damping factor 102
DC-input modules 366
DC-output modules 367
DC-resonant charging 89
DC-to-dc converters 7, 311
Delta modulation 331
Delta transformer connection 263
"Delta-wye" transformer configuration 269
Deuterium 128
Dielectric oil 39
Diode (e-gun). *See* cathode-controlled electron gun
"Diode line" 167
Doppler ambiguity 119, 249
Doppler sidelobe 118, 119
Droop 149
Dummy load 28
Duplexer
 antenna. *See* Antenna duplexer
 balanced 24
 branch 24

Duty cycle. *See* Duty factor
Duty factor 5, 46

E

e-gun. *See* electron gun
EIMAC 66
Electric-field gradient 57
Electrical interlock switches 34
Electrical length 113
Electromagnetic compatibility 33
Electromagnetic susceptibility 32
Electron gun 65
Electronic crowbar 152, 253
 definition 6
Electronic voltage regulator
 definition 6
ELF. *See* extremely low frequency; extremely-low-frequency magnetic fields
ELF submarine communication system 319
Emission density 165
"End-of-pulse switch-tube arc" 196
Energy absorption
 resistance 261
"English-muffin" effect. *See* "microwave oven" effect
Ergonomics 40
Exciter 16
Expanded-pulse input 108
Exponentially decaying waveform 50
"Extended-delta" transformer configuration 271
Extremely low-frequency magnetic fields 39
Extremely-low-frequency (ELF) 319

F

Failure modes and effects analysis 40
Fault current limiting
 definition 6
Fiber-optic coupling 132
Field emission 123
Filament supply 4
Film coefficient 339
Film conductance 339
Filter
 "leaky wall" 24
 "waffle"-type 24
"Flicker" modulation 279
Flip-flop 363
"Floating deck" 152
Flow-interlock electrical switches 345
Flyback converter 316
Focus electrodes 4
Focus magnet supply 4
"Forced" commutation 320
Forced convection 339
Four-layer device 131, 134
Fourier series 264
Free-electron lasers 113

Free-wheeling diode 311
 Frequency-compensated voltage divider 371
 Full-control switch 148

G

Gate-catcher 208
 Global resistor 262
 "Graceful degradation" 32
 Gradient grid 123
 Grid bias 148
 Grid blackening 166
 Grid firing 124
 Grid-catcher 206
 "Grid-leak" bias 238
 Ground-bounce voltage 377
 Guillemin networks 98
 Type A 98
 Type B 99
 Type C 99
 Type D 99
 Type E 99
 Type F 100

H

Half-bridge inverter 317
 Half-control switch 121
 Harmonic components 277
 Harmonic emanations 33
 Harmonic filter 23
 Harmonic pollution 277
 Harmonic resonance 23
 HAX. *See* Haystack Auxiliary Radar
 Haystack Auxiliary Radar 232
 Heat transfer 334
 Heat-removal rate 343
 High- μ triode 166
 High-frequency electronic power conditioner 329
 High-frequency three-phase alternators 322
 High-noise-immunity logic (HINIL) 363
 High-power amplifier 20
 High-speed contactor 8, 36, 306
 High-threshold-logic (HTL) 363
 High-voltage cable insulation 63
 High-voltage transformer/rectifier 7, 263
 HPA. *See* high-power amplifier
 Hydrogen 122
 Hydrogen thyratron 90, 121
 Hydrogen-reservoir heater 123
 Hypervapotron 351

I

Ignitor 129
 Ignitron 40, 129, 142
 Incremental conductance 113

Incremental efficiency 113
 Induction voltage regulator 292
 "Inductive-kick" 247
 "Inductrol" 292
 "Infinite-voltage-range" crowbar 259
 Injectron beam-switch tube 184
 Interlocks. *See* electrical interlock switches
 Interrupting rating 9
 Intrapulse droop 153
 Ionization 124
 Ionizing radiation 39
 Isolator
 ferrite 18
 IVR. *See* induction voltage regulator

K

"Kirk" key system 38
 Klystron 2-4

L

L-5012 Injectron 184
 L-5097 Injectron 184
 L-5773 klystron 222
 Large-orbit gyrotron 94
 Laser trigger 133
 Leakage inductance 102
 transformer 9
 Leidenfrost point 350
 Line-type pulser 71
 Line-type pulser discharge switches 121
 Linear phase change 150
 Linear-chirp frequency modulation 106
 Load lines 189
 Long-range imaging radar 229
 LRIR. *See* Long-range imaging radar

M

Magnetic assist 125
 Magnetic force 101
 Magnetic modulator 132
 Magnetic pulse compression 132
 Magnetizing inductance 102
 Maintainability 31
 Manual load-break disconnect 308
 MAPS-40 thyratron 127
 MAR-1. *See* Multi-Function Array Radar 1
 Marx impulse generator 97
 Mean time between failure 29
 Metal-oxide semiconductor field-effect transistor. *See* MOSFET
 Metal-oxide varistor 154
 "Metallic" noise 278
 Micro-projections 255
 "Microwave-oven" effect 39
 Mid-plane electrode 258

- MIL-HDBK 217 29, 30, 31
"Miller" capacitance 193
Millstone Hill radar transmitters 221
"Mismatch loss" 79
ML- 6544 triode 170
ML- 8618 triode 184
ML-8549 triode 179
Modulating anode 4, 66
Modulating-anode pulsers 200
Modulation sideband 111
Modulator gate 243
MOSFET 30, 193
MOV. *See* metal-oxide varistor
Moving-target indicator 118
MTBF. *See* mean time between failure
MTI. *See* moving-target indicator
Multi-function Array Radar 1, 226
- N**
- NAND gate 362
Needle-gap geometry 60
Nucleate boiling 350
Nusselt number 340
- O**
- Oil figure of merit 60
Oil voltage-breakdown test 60
"Open-delta" transformer connection 286
Optically coupled links 152
Output autotransformer 144
Output-pulse transformers 101
Overmatch condition 78
Overshoot 153
Oxide-coated cathodes 162
- P**
- PARCS. *See* Perimeter Acquisition Radar Characterization System
Partial discharge 155. *Also see* corona
Paschen's curve 54
PB factor 127
Perimeter Acquisition Radar Characterization System 32
Permanent magnet 227
Perveance 147
PFN. *See* pulse-forming network
Phase deviation 109
Phase modulation 110
Phase shifter 19
Phase-angle "pulling" 112
Phase-controlled rectifier 282
Phase-delay 112
Phase-pushing factor 112, 114
Phased-array transmitter 12
PIN diode switch 16
PIN diodes 25
Plasmoid trigger 136
PLC 10. *See also* programmable logic control; programmable-logic controller
PM. *See* phase modulation
"Polygon" transformer configuration 271
Polyphase circuits 263
Power contactors 306
Power factor 274
Power grid vacuum tubes 161
Power splitter 18
PRF. *See* pulse-repetition frequency
PRI. *See* pulse-repetition interval
Primary SCR controller 284
Programmable logic control 10, 365
"Projected" plate-current cutoff 168
Pseudo-spark thyatron 128
Pulse "balun" 376
Pulse compressor 110
Pulse modulator 5
Pulse stretching 166
Pulse-forming network 75, 98
 charge regulation 120
Pulse-repetition frequency 46
Pulse-repetition interval 46
Pulse-top flatness 106
Pulse-top ripple 107
Pulse-width-modulated regulator 312
Pulsed energy storage 5
Purity-maintenance loop 345
Push-pull converter 316
PWM. *See* pulse-width-modulated regulator
PWM regulator 329
Pyrolytic graphite 166
Pyrovac 335
- Q**
- Quasi-resonant inverter 326
"Quasi-resonant" modulating-anode pulser 210
- R**
- "Rabbit-ears" 237
Radial-beam tube 146
Radiation 334
Rayleigh line 137
RBDT. *See* reverse-blocking diode thyristor
RBTT. *See* reverse-blocking triode thyristor
Relativistic klystron 94
Relay-logic ladder-network 352
Reliability 29
Reverse-blocking diode thyristor 134
Reverse-blocking triode thyristor 134
Reverse-switching rectifier (RSR) 134
Reynolds number 340
RF breakdown 12
RF driver amplifier 18
"Ribwound" resistor 261

Ring-type current-monitor transformer 370
 Ringing 153
 RMS. *See* root-mean-square current
 Root-mean-square current 47
 Rotameter 345

S

S94000E tetrode 173
 Safety 3s
 "Scalloping" 234
 SCR 131. *Also see* thyristor
 SCR primary controller 8
 variable voltage device 8
 Secondary emission 166
 Selectable-contact output module 367
 Self-capacitance of modulator decks 215
 "Serrodyne" effect 151
 SF₆. *See* sulfur hexafluoride
 Shadow grid 234
 Shoot-through 318
 Short-circuit current 52
 Shunt-trip circuit breaker 36, 308
 Silicon controlled rectifier. *See* SCR
 Single-quadrant forward converter 315
 Single-sideband modulation 109
 Six-phase, full-wave (12-pulse) rectifier 268
 SLAC. *See* Stanford Linear Accelerator Center
 "Soft-tube" 75
 "Soft-tube" modulator 71
 Solenoids 4
 Solid dielectrics 61
 Space-charge-limited electron flow 149
 Spark gaps 134
 Spike leakage 25
 Spurious emanations. *See* harmonic emanations
 Stanford Linear Accelerator Center 136
 Step regulator 291
 Step start (or soft start) 7
 Storage capacitors 155
 Sulfur hexafluoride 39
 Surge-limiting resistors 261
 Switch-mode electronic power conditioning 311
 Switch-mode power-conditioning 140
 Switchgear 306
 Switching loss 90
 "Synchronous vibrator" 320

T

Technical proposal 43
 Technical specification 43
 Telephone Influence Factor (TIF) 278
 Temperature rise 344
 Tetrode 146
 TH-5184 tetrode 220
 TH-5188 tetrode 220
 Thermal management 334

Thermal path 336
 Thermal resistance 337
 Thermal resistivity 337
 Thermosyphoning 351
 Thoriated tungsten 162
 "Three-bump" rectifier. *See* three-pulse rectifier
 Three-phase, full-wave (six-pulse) rectifier 267
 Three-pulse rectifier 265
 Thyristor 131
 Thyrite 154
 Time-delay shim 19
 "Time-sidelobe" 108
 "Tin-foil" test 255
 "Totem-pole" inverter. *See* half-bridge inverter
 Transactor 120
 Transducer 120
 Transformer oil 59
 Transistor-transistor logic 363
 Trapped air 61
 "Trapped" resonances 277
 Traveling-wave-tube amplifier 329
 Trigger electrode 259
 Triggered vacuum gap 135
 Triode 146
 TTL. *See* transistor-transistor logic
 Turbine-type flow meter 347
 Turbulence 340
 TVG. *See* triggered vacuum gap
 TWTA. *See* traveling-wave-tube amplifier

U

Under-voltage trip
 circuit breaker 36
 Undermatched case 81
 Unit substation 309
 definition 8

V

Vacuum contactors 306
 Vacuum tube 146
 Vapor-phase cooling 347
 Variable auto-transformer 8, 288
 Variable voltage device 8, 282
 definition 7
 Variable-voltage transformer 289
 "Variac" 288
 Velocity-jump guns 200
 Vibrator 320
 "Virtual cathode" 162
 Virtual-cathode oscillator 94
 Volt-ampere demand
 unit substation 9
 Voltage comparator 375
 Voltage droop 102
 Voltage reversal 155

Voltage sample 6
Voltage-multiplier rectifier 322
VTU-5692 TWT 232
VTX-5681 TWT 229
VVT. *See* variable-voltage transformer

W

Waster load 19, 23
Water-filled Blumlein 95
Waveguide 12
 beam 14

quasi-optical 14
 switch 28
WL-8461 triode 170

X

X-780 klystron 221

Y

Y-540 245
Y-847 triode 218
YU-146 tetrode 219

LOS ALAMOS NAT'L LAB.
IS-4 REPORT SECTION
RECEIVED

'94 JUN 13 PM 1 33

This report has been reproduced directly from the best available copy.

It is available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831. Prices are available from (615) 576-8401.

It is available to the public from the National Technical Information Service, US Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.



Los Alamos
NATIONAL LABORATORY

Los Alamos, New Mexico 87545