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

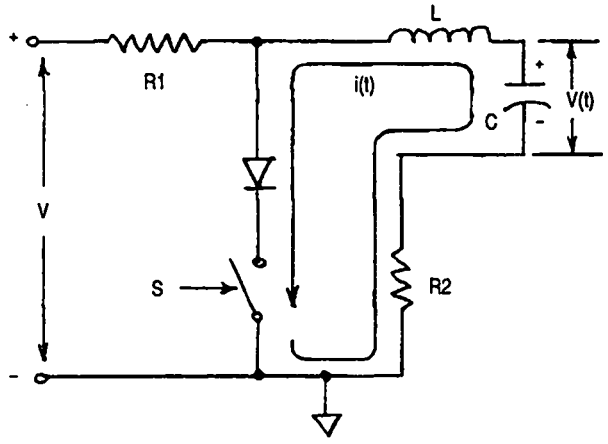

Figure 9-1. The simplest form of a line-type modulator (R-L-C circuit).

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, $R2$. 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 $R2 = R = 0$, 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

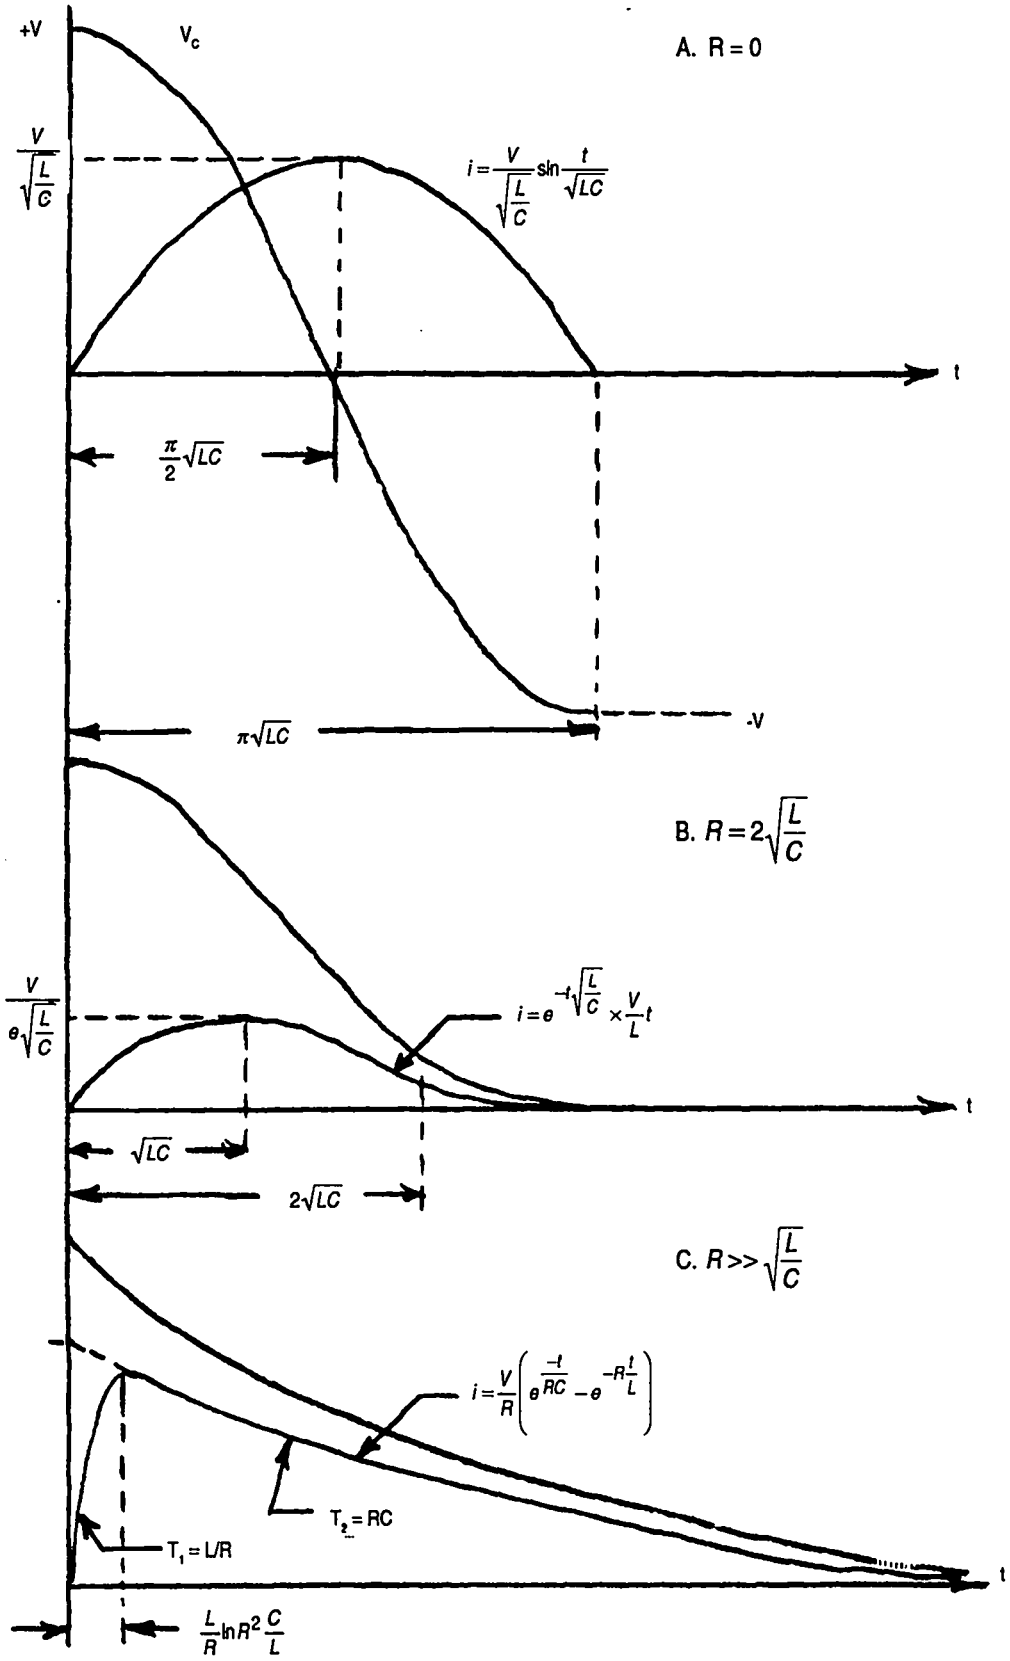

Figure 9-2. Discharge characteristics of a simple R-L-C circuit.

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 $2\sqrt{(L/C)}$, 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:

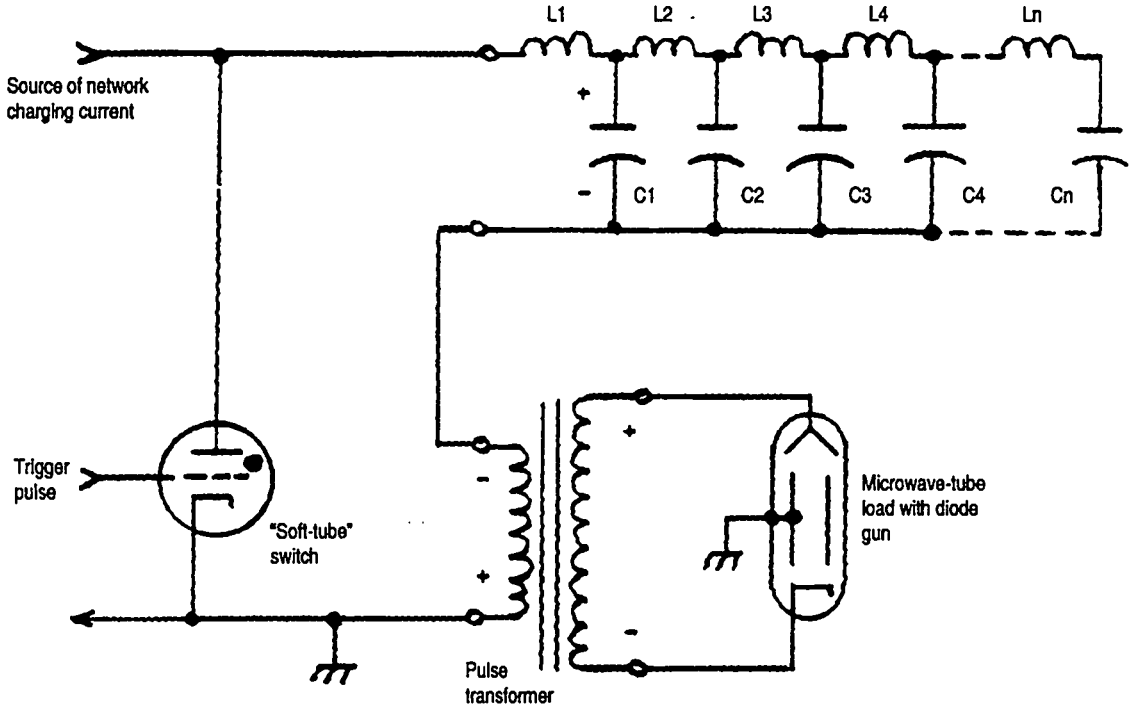

Figure 9-3. Simplified discharge circuit of pulser with artificial transmission-line pulse-forming network.

$$\text{Load energy} = W_L = \int_0^\infty i^2 R dt = \int_0^\infty e^{\frac{-2t}{\sqrt{LC}}} \left(\frac{V}{L}t\right)^2 \times 2\sqrt{\frac{L}{C}}dt$$

$$= \left[\frac{2V^2}{L^2}\sqrt{\frac{L}{C}} \frac{e^{\frac{-2t}{\sqrt{LC}}}}{8LC\sqrt{LC}} \left(\frac{4t^2}{LC} - \frac{4t}{\sqrt{LC}} + 2\right) \right]_0^\infty$$

$$= \frac{2V^2}{L^2}\sqrt{\frac{L}{C}} \times \frac{2LC\sqrt{LC}}{8} = \frac{4}{8}CV^2 = \frac{1}{2}CV^2 .$$

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 x 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^2C/L after the switch closure. The peak value of current is slightly less than V/R, or

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

(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

$$W_0 = \frac{1}{2} C V_0^2 .$$

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

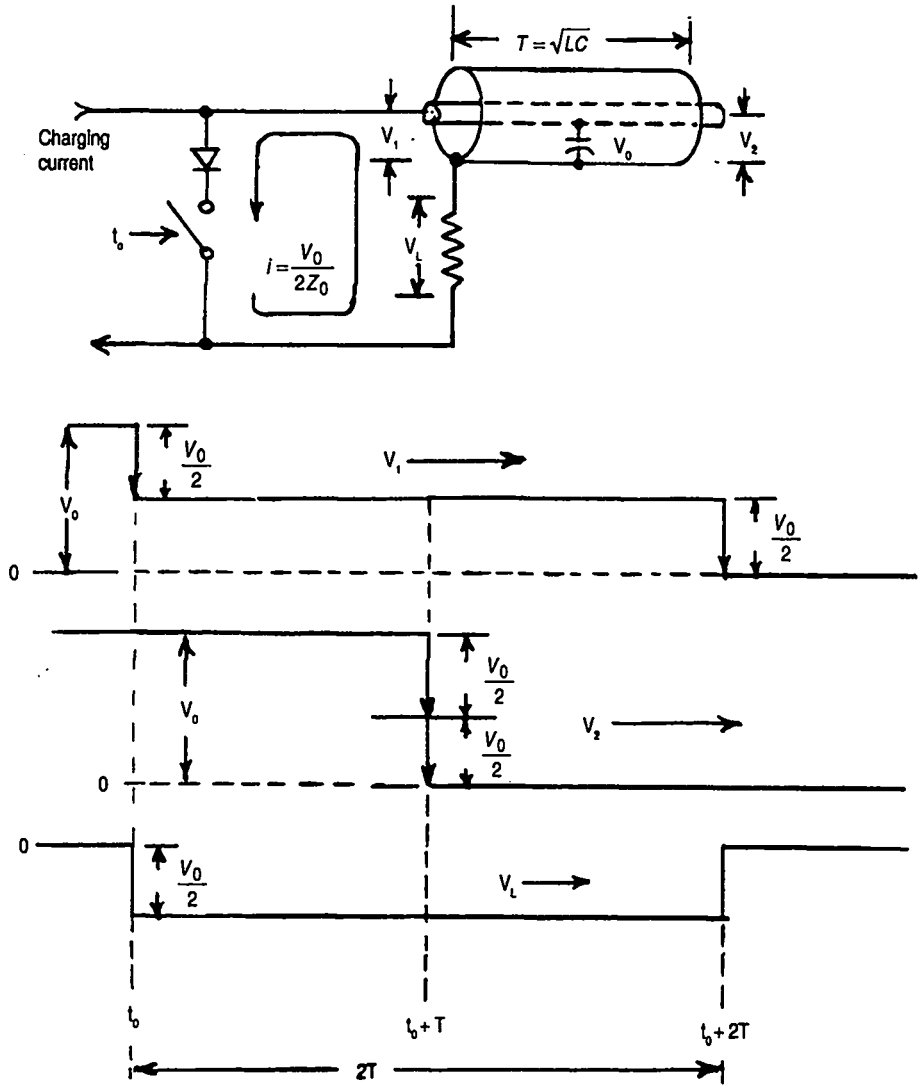

Figure 9-4. Discharge of transmission-line pulse-forming network into matched load.

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:

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

$$L = CZ_0^2$$

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

$$W_L = \frac{1}{2}Li^2 = \frac{1}{2}CZ_0^2 \times \left(\frac{V_0}{2z_0}\right)^2 = \frac{CZ_0^2V_0^2}{8Z_0^2} = \frac{1}{8}CV_0^2 = \frac{1}{4}W_0.$$

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

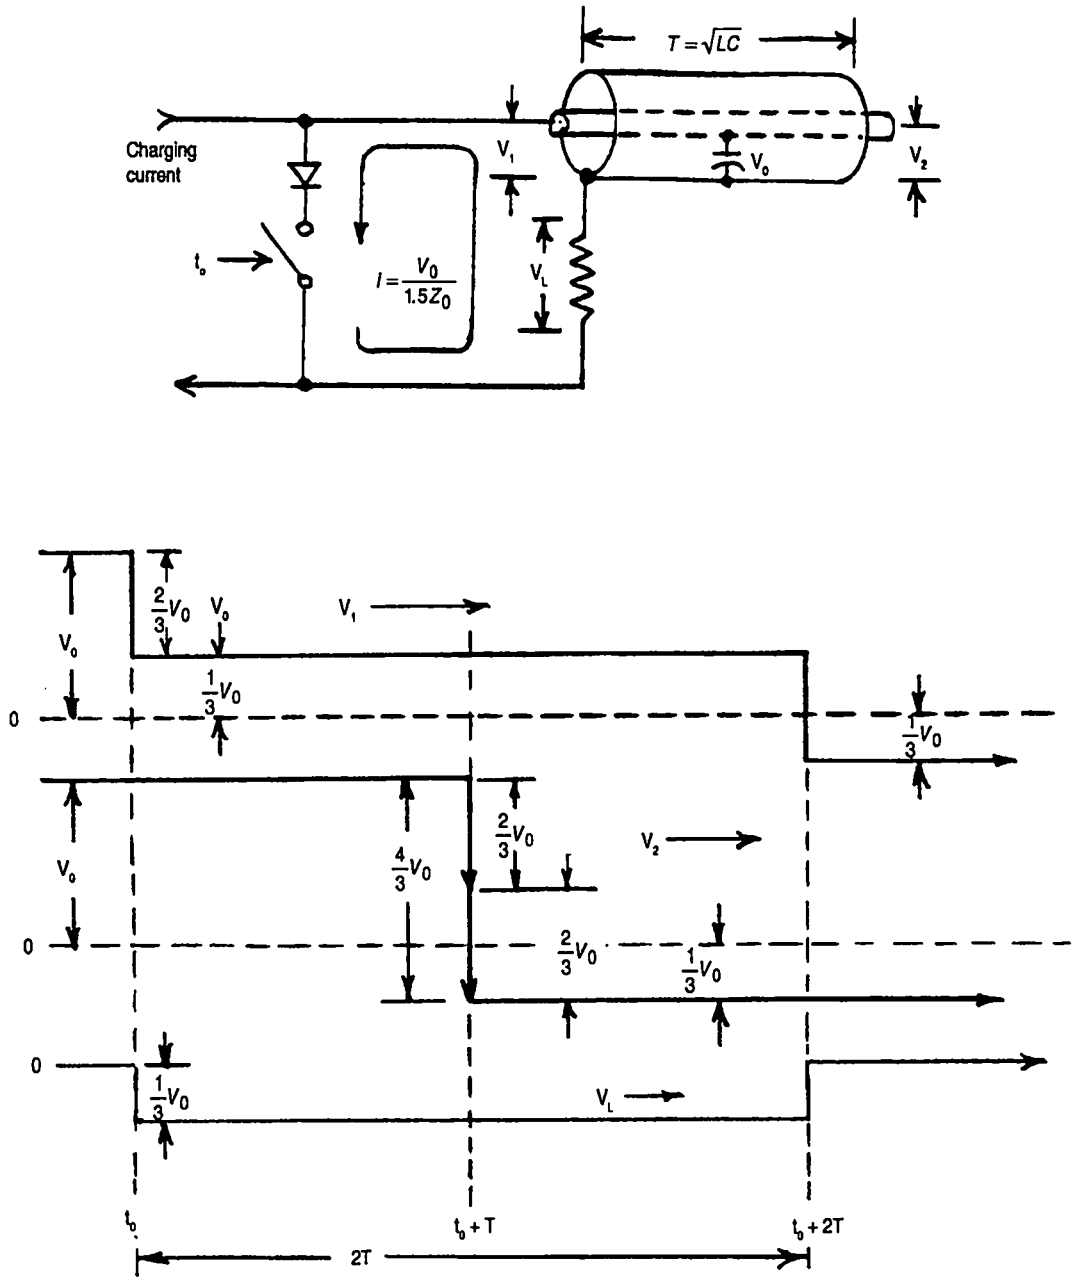

Figure 9-5. Discharge of transmission-line pulse-forming network into 2:1 overmatch.

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

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:

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

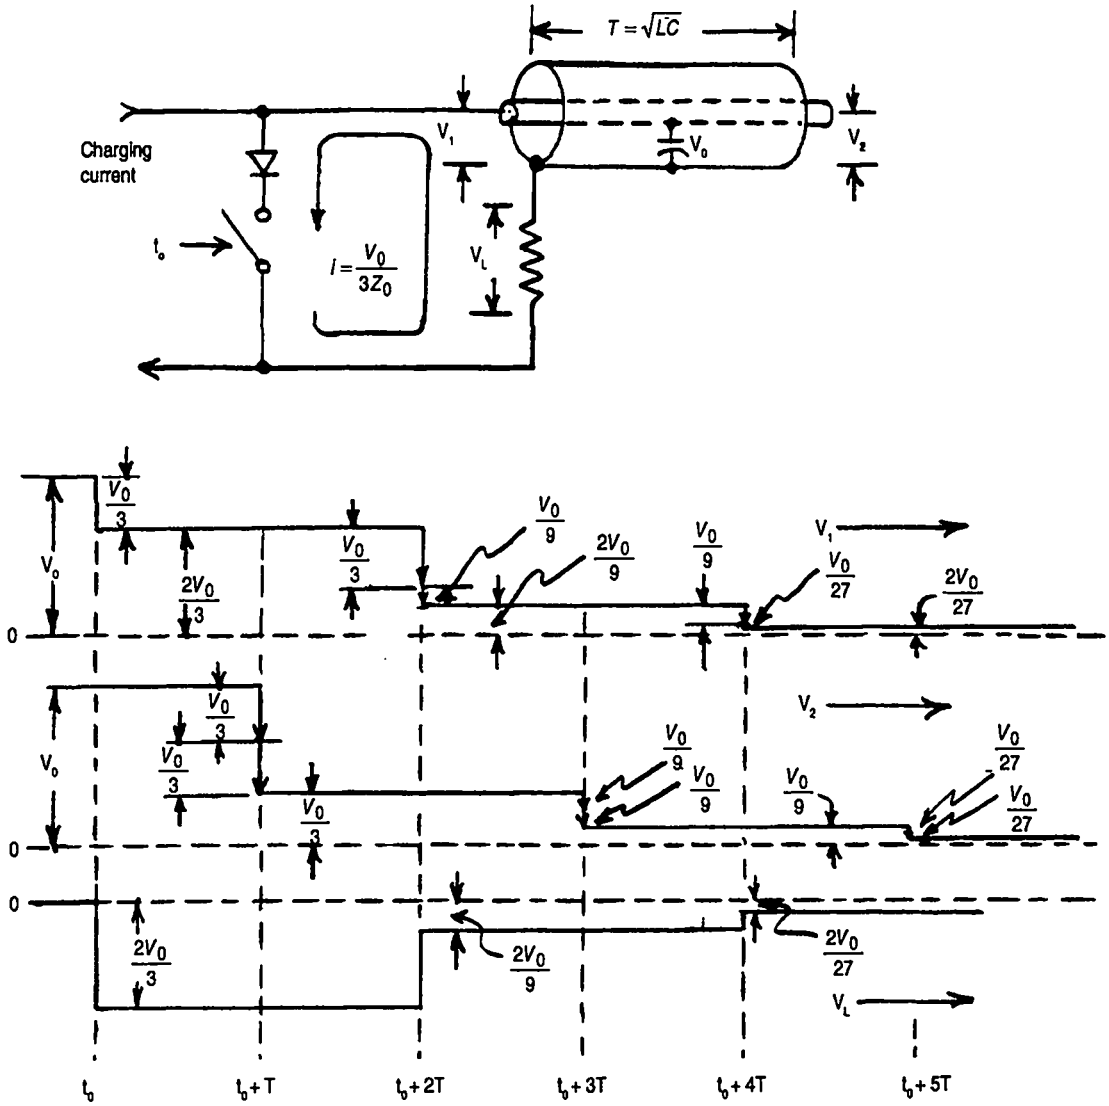

Figure 9-6. Discharge of transmission-line pulse-forming network into 2:1 undermatch.

$$W_{stored} = \frac{1}{2}C\frac{(V_0)^2}{3} = \frac{1}{9}\times\frac{1}{2}CV_0^2 = \frac{1}{9}W_0.$$

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

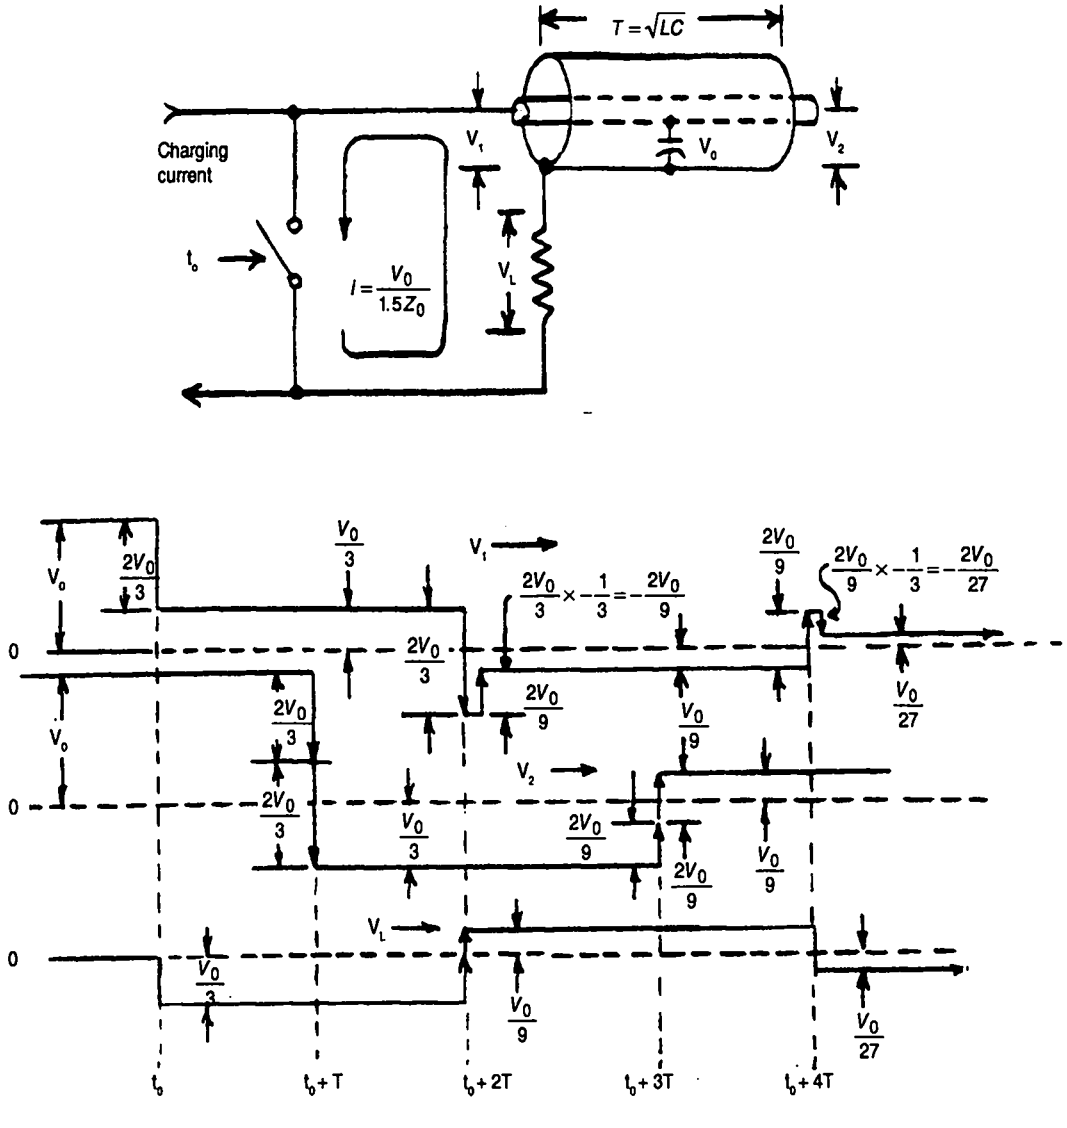

Figure 9-7. Discharge of 2:1 overmatched network with bidirectional switch.

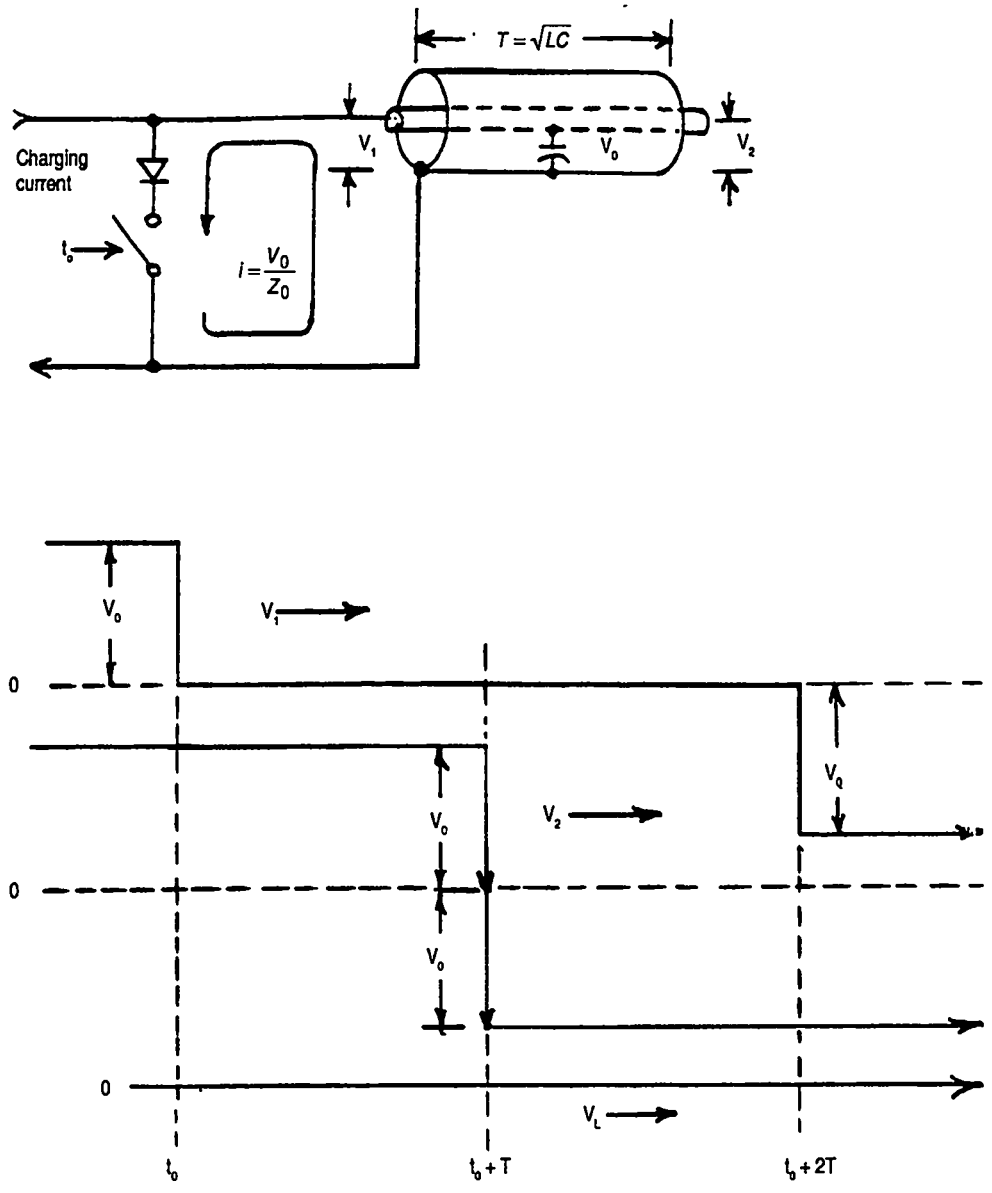

Figure 9-8. Discharge of transmission-line network into short-circuit load.

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.

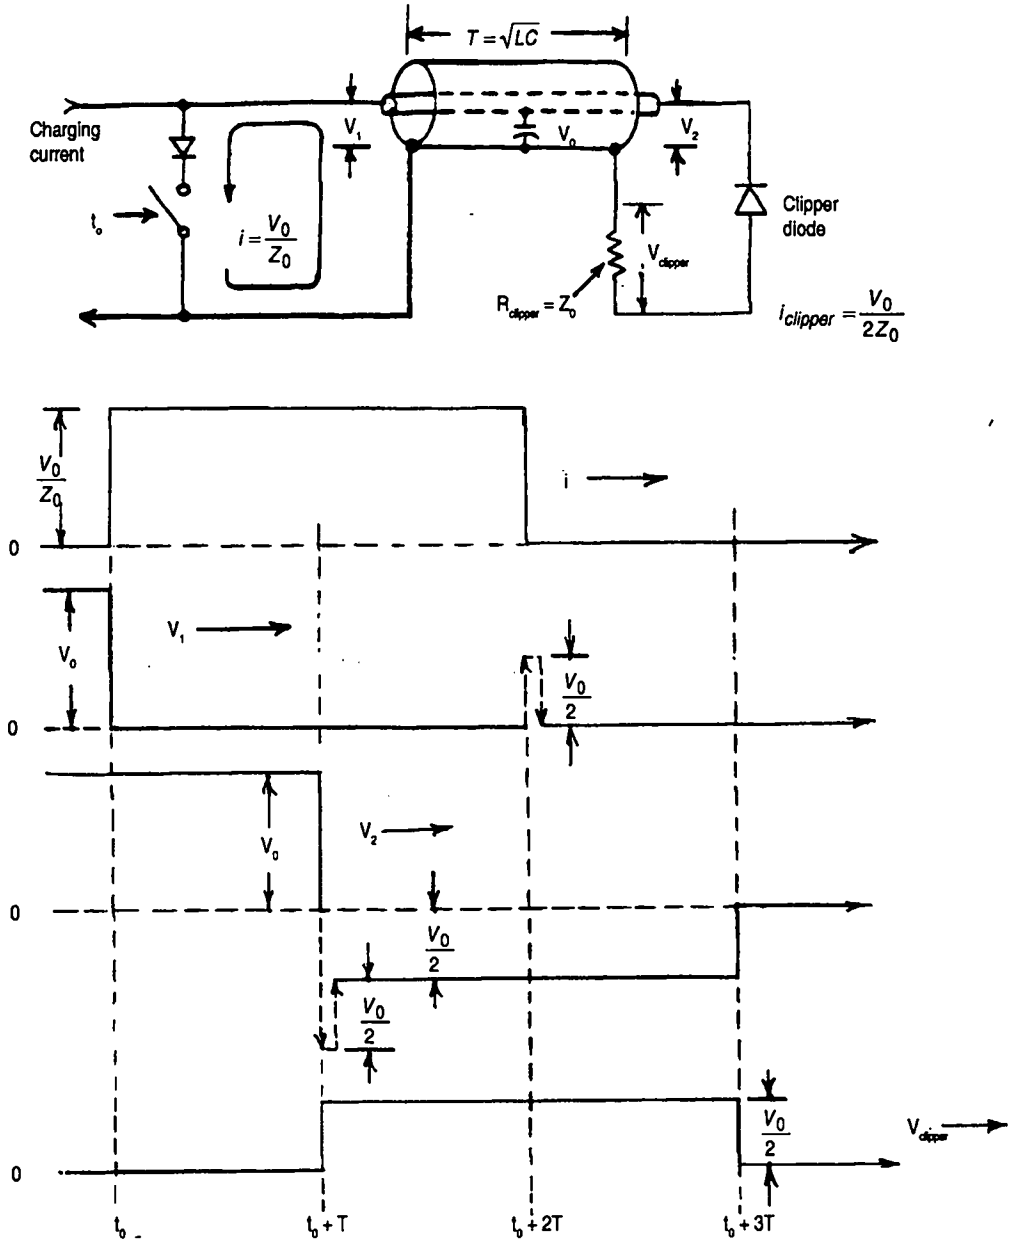

Figure 9-9. Discharge of transmission-line network with end-of-line clipper circuit into short-circuited load.

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

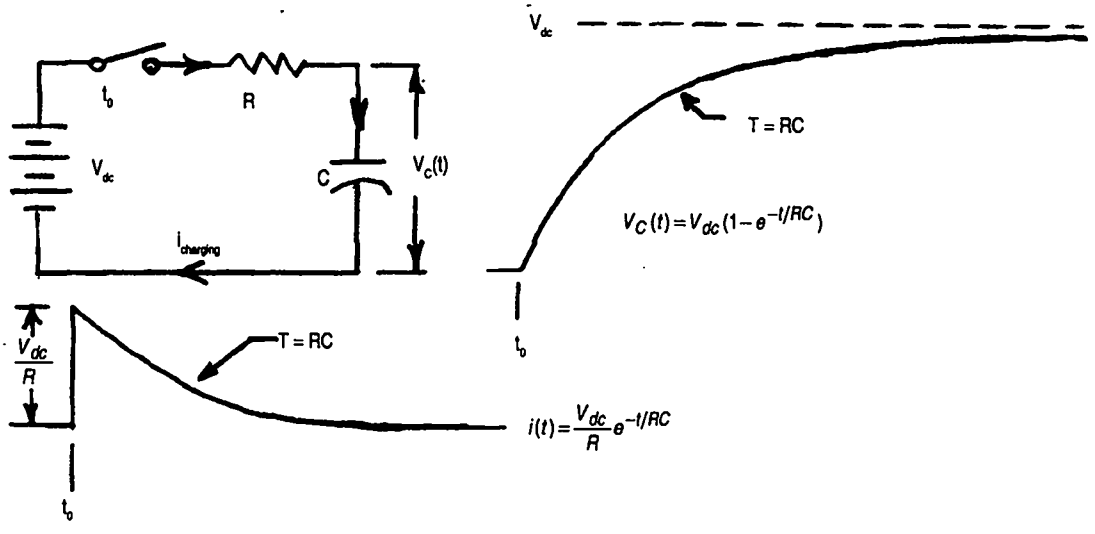

Figure 9-10. Capacitor charging with resistive current limiting.

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,

$$W_c = \int_0^\infty i(t) \times V_c(t)dt$$

$$= \int_0^\infty \frac{V_{dc}}{R} e^{-t/RC} \times V_{dc}\left(1 - e^{-t/RC}\right)dt = \frac{V_{dc}^2}{R}\int_0^\infty \left(e^{-2T/RC}\right)dt$$

$$= \frac{V_{dc}^2}{R}\left[_1^\infty - RCe^{-t/RC} + \frac{RC}{2}e^{-2t/RC} = \frac{V_{dc}^2}{R}\left(RC - \frac{RC}{2}\right) = \frac{V_{dc}^2}{R} \times \frac{RC}{2}\right.$$

$$= \frac{1}{2}CV_{dc}^2,$$

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

$$V_{dc}\int_0^\infty i_{charging}(t)dt = V_{dc}\int_0^\infty \frac{V_{dc}}{R}e^{-t/RC}dt = \frac{V_{dc}^2}{R}\left[_0^\infty - RCe^{-t/RC}\right.$$

$$= \frac{V_{dc}^2}{R} \times RC$$

$$= CV_{dc}^2,$$

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

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

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

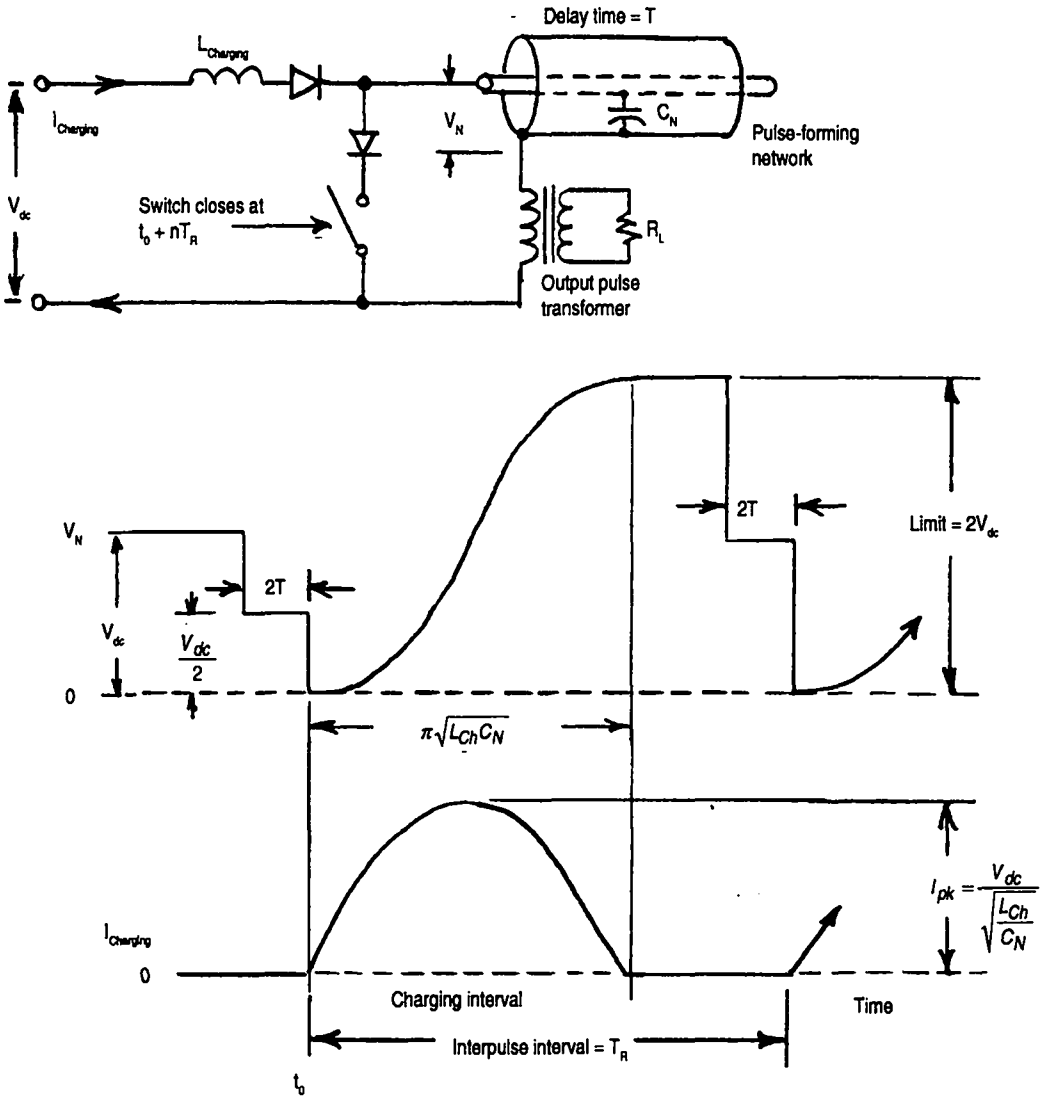

Figure 9-11. DC-resonant charging of pulse-forming network.

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

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:

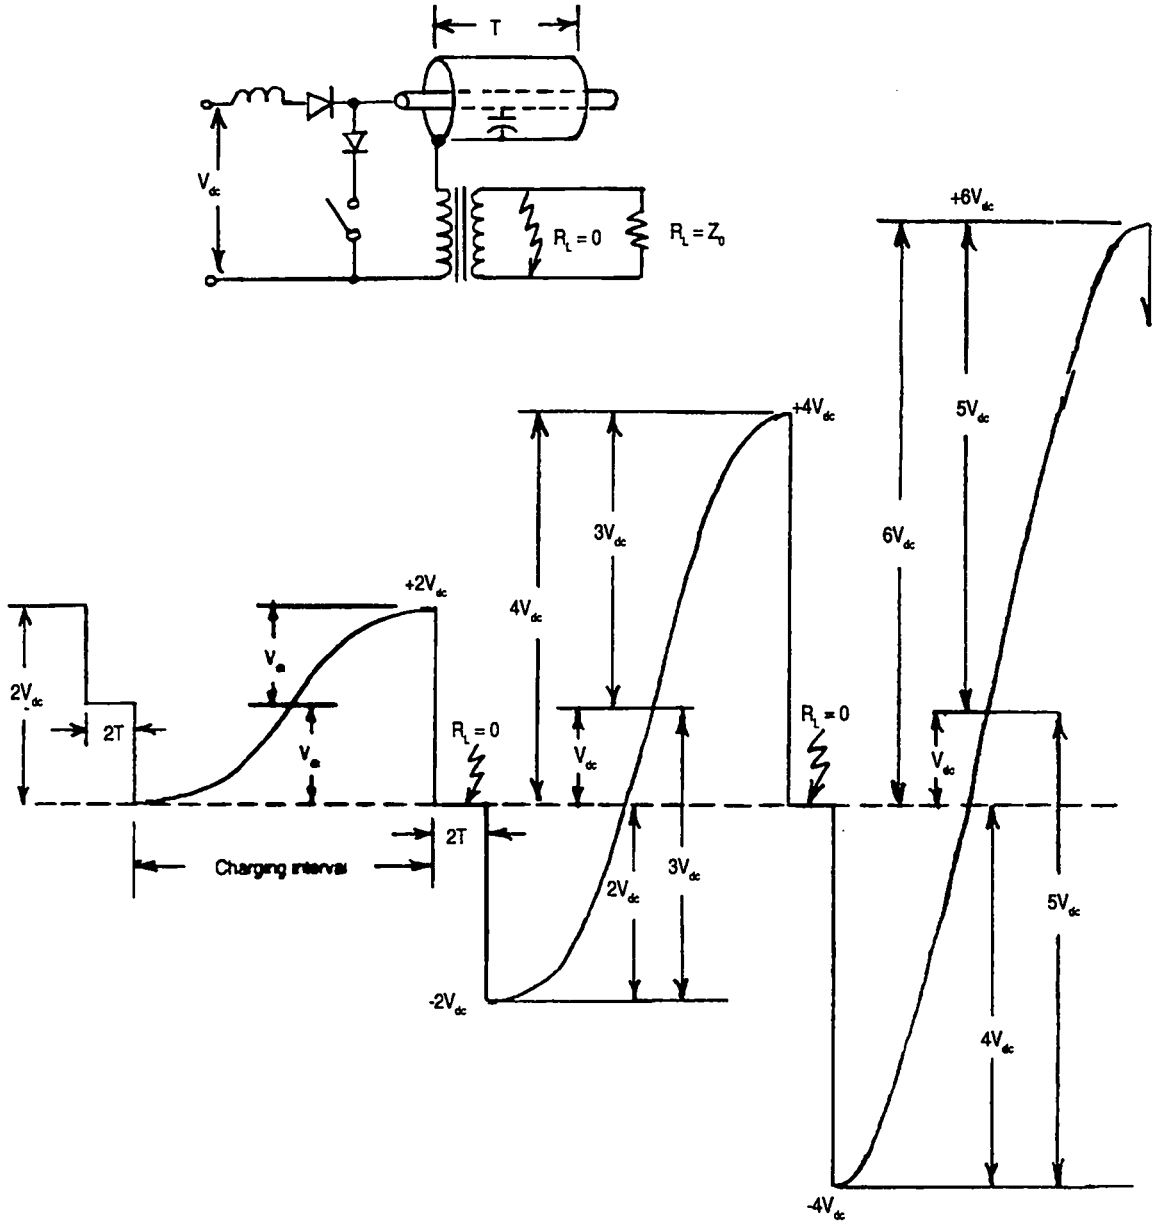

Figure 9-12. DC-resonant charging with shorted load and no clipper circuit.

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

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

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 $-2Vd_c$ 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 ohms/100 = 10 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 thyratron. 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 kV/(2 x 10 ohms), or 2500 A. The average thyratron 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 thyratron, 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 thyratrons 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 x 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

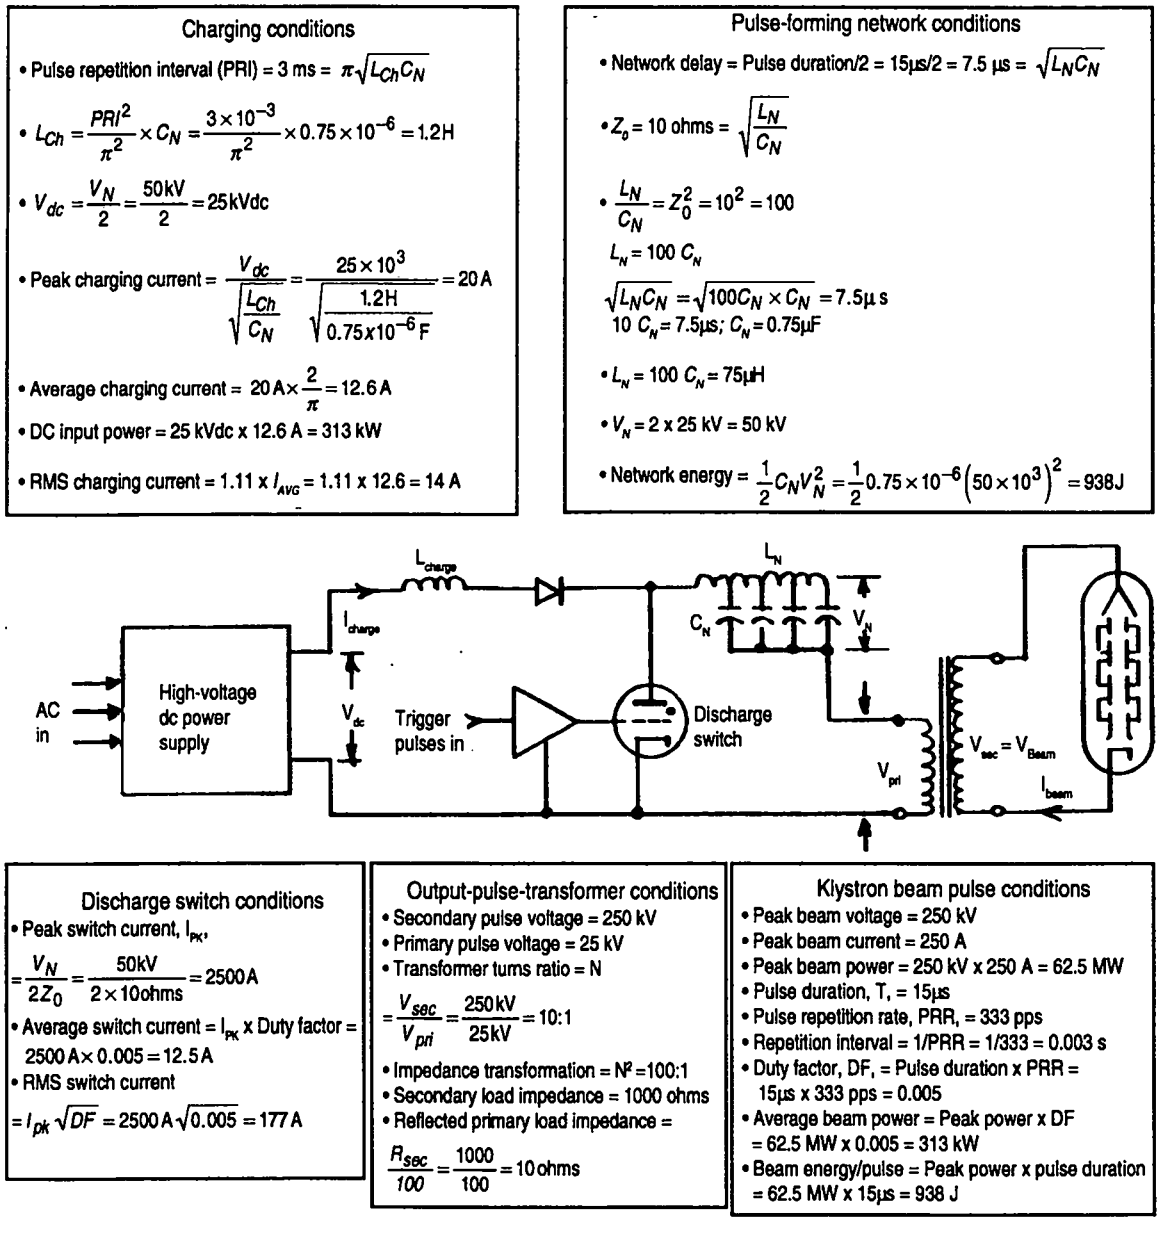

Figure 9-13. An illustrative example of a high-power, line-type modulator.

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 thyratron, 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 thyratron 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.

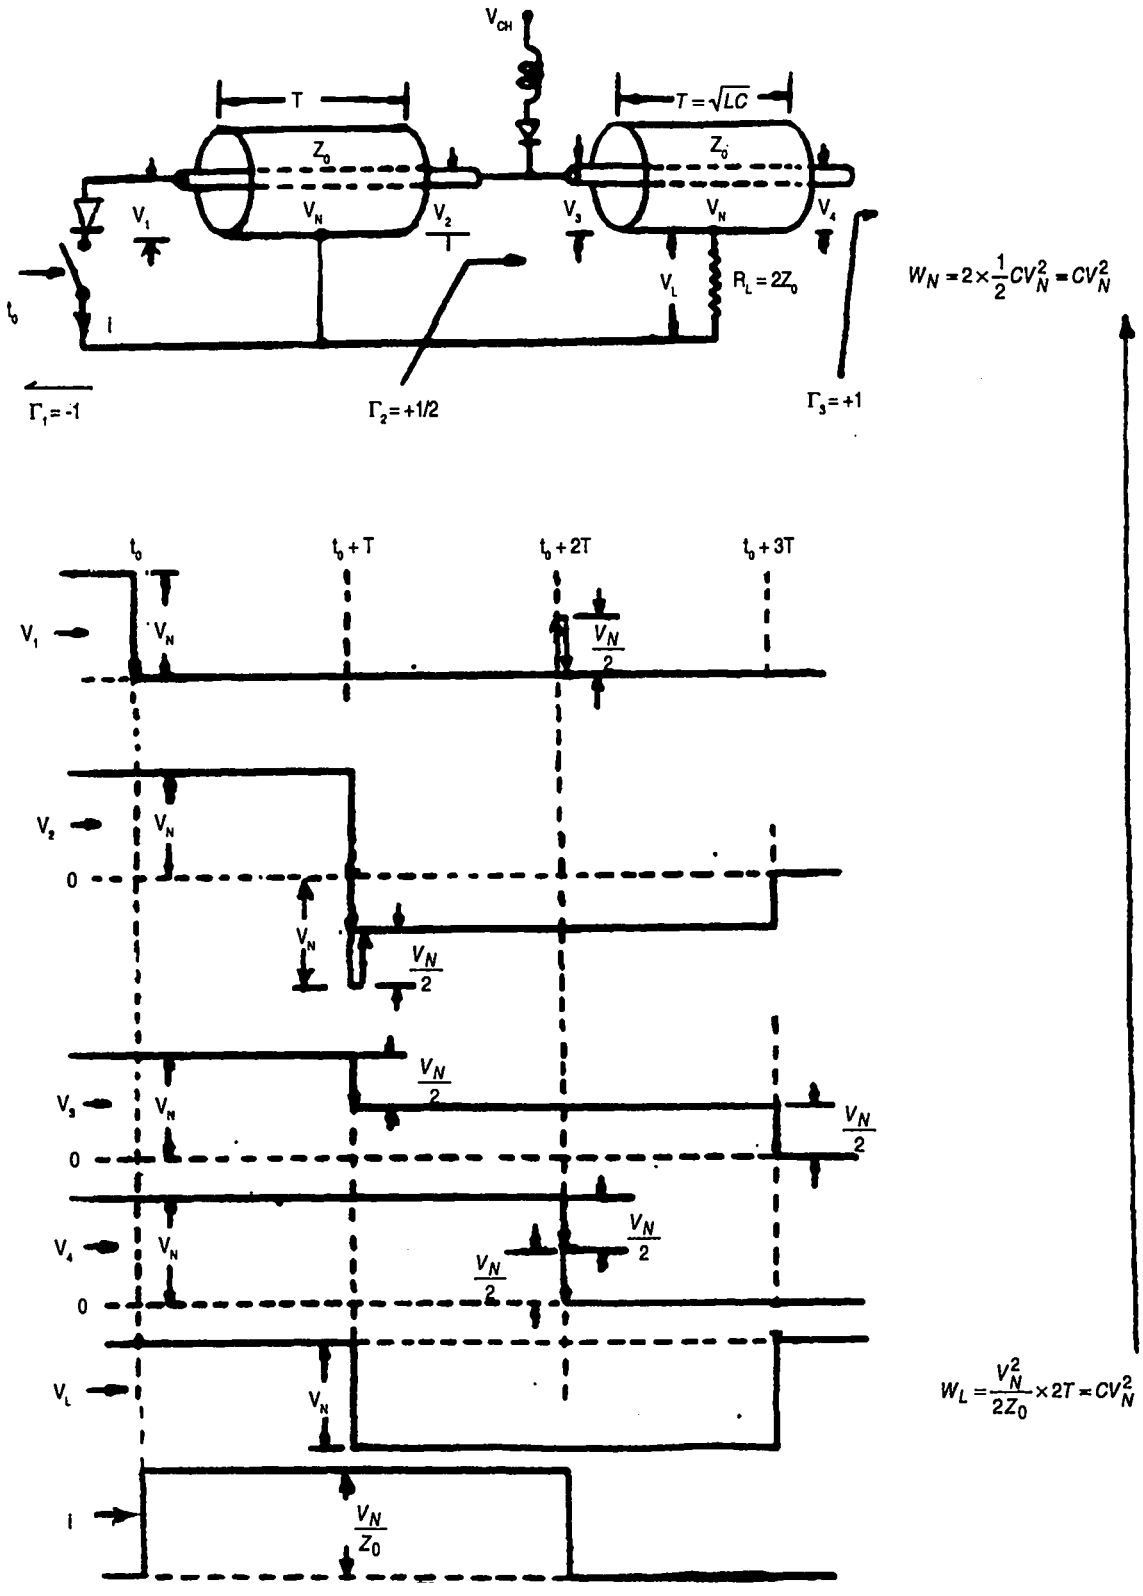

Figure 9-14. Discharge of the matched-load Blumlein circuit.

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_NC, 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

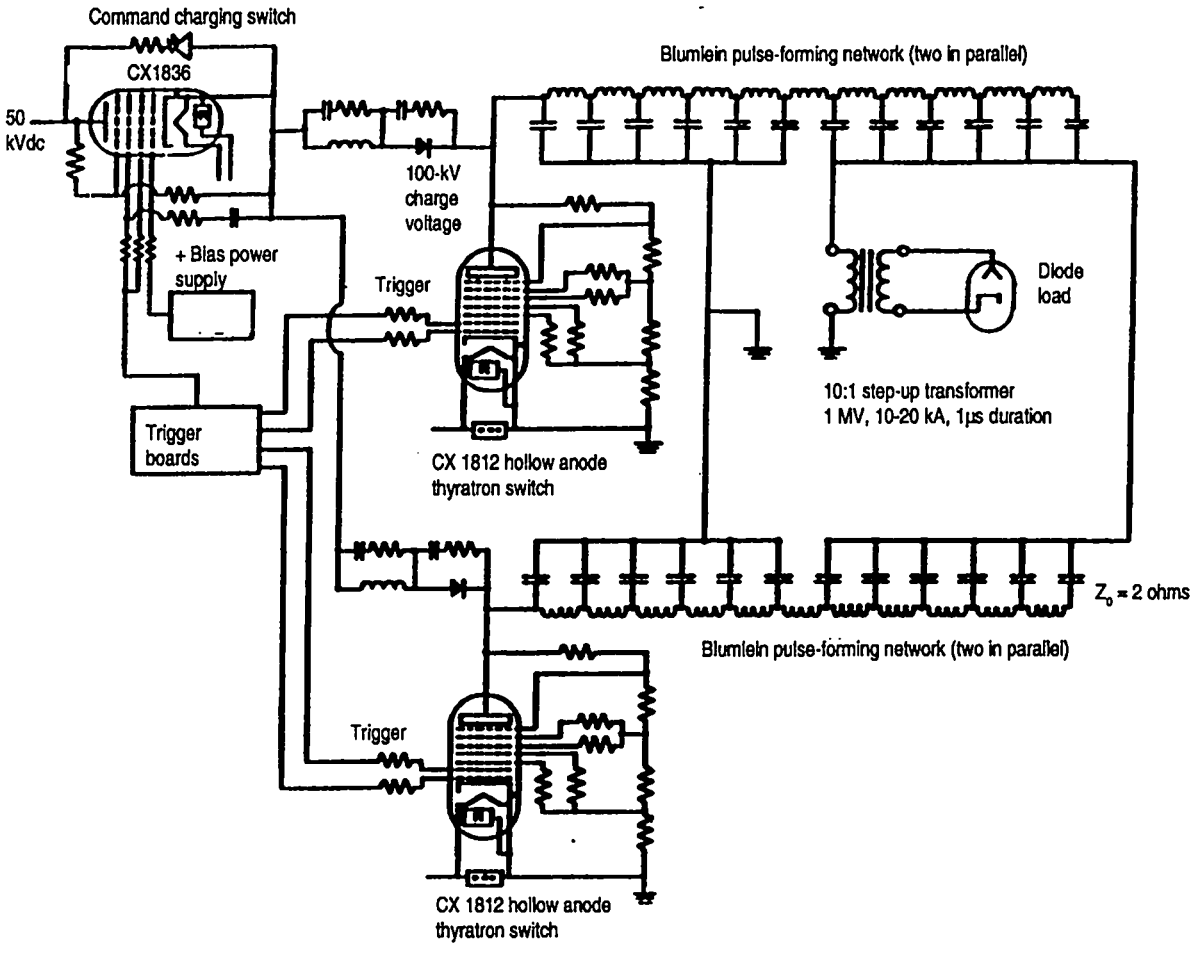

Figure 9-15. The BANSHEE, a megavolt-output Blumlein-type pulse modulator.

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

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

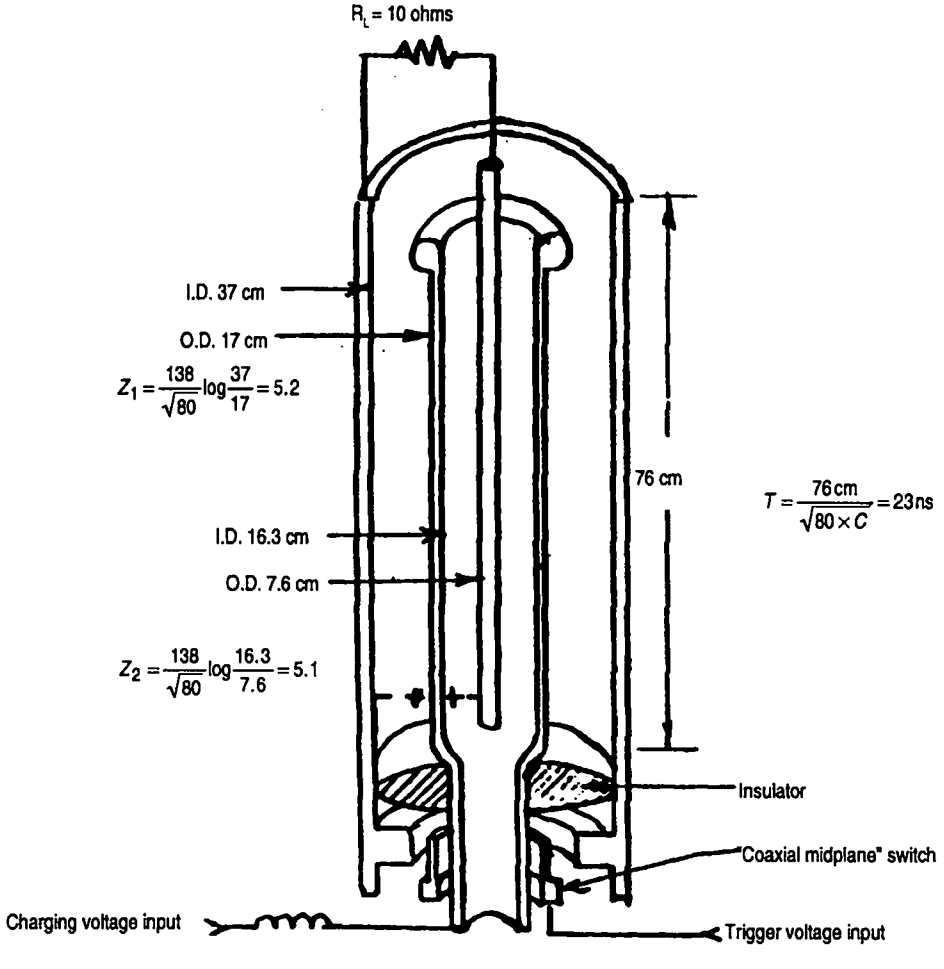

Figure 9-16. The water-filled coaxial short-pulse, high-power Blumlein network.

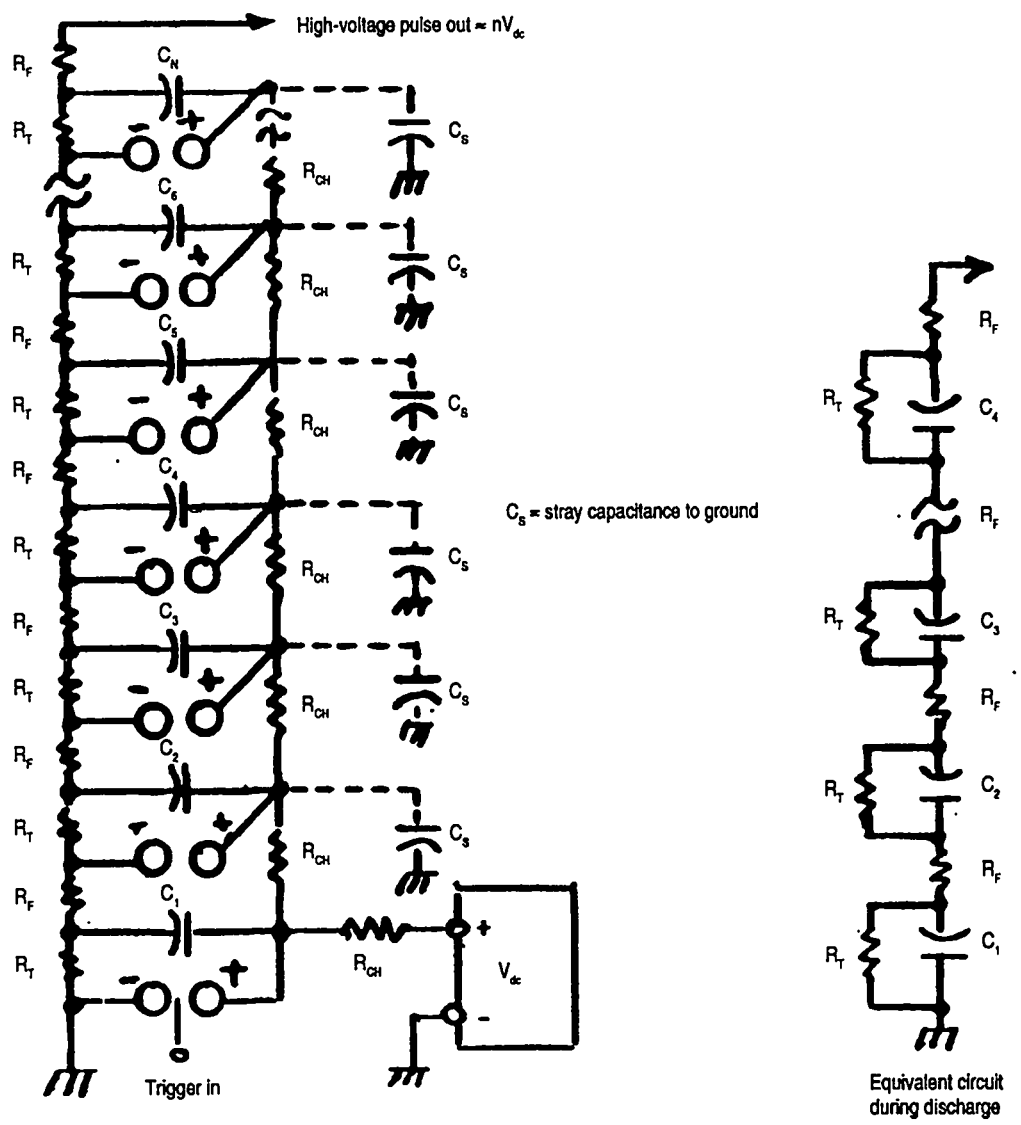

Figure 9-17. Schematic representation of the Marx-type impulse generator.

shown, the characteristic impedances of both lines were approximately 5 ohms. When filled with high-purity water, the dc capacitance was 0.01 μF 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 μs—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 resistance 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.

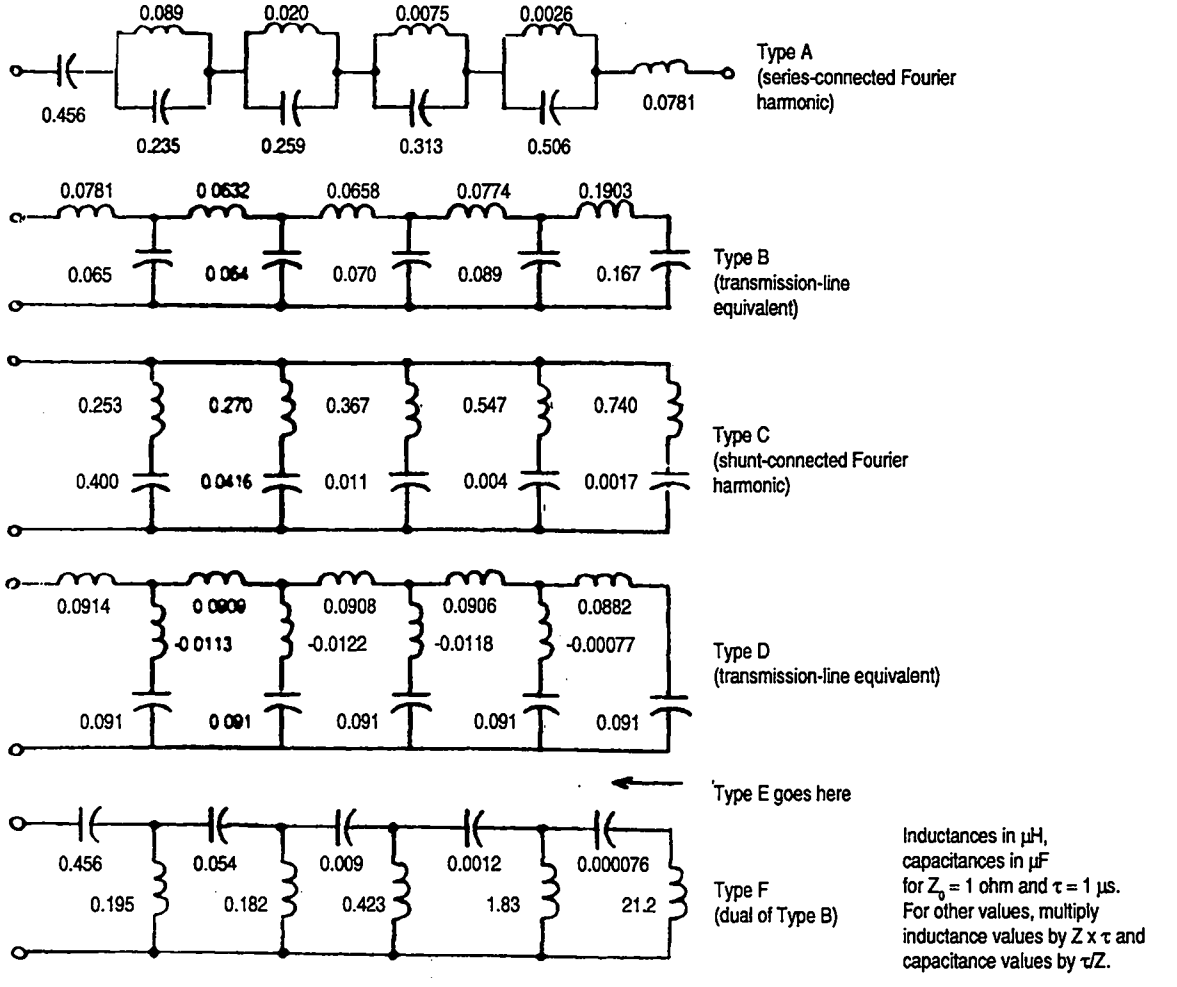

Figure 9-18. The Guillemin lumped-circuit pulse-forming network types.

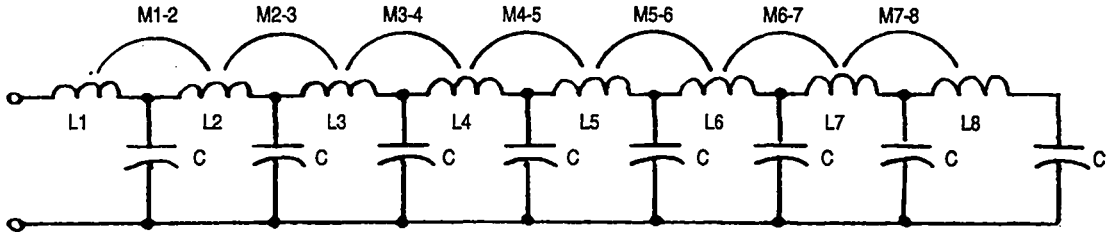

Number of sections	C	M1-2	M2-3	M3-4	M4-5	M5-6	M6-7	M7-8
		L1	L2	L3	L4	L5	L6	L7	L8
4	0.125	0.195	0.00829	0.0973	0.0144	0.0953	0.0131	0.107	
6	0.0838	0.140	0.00490	0.0657	0.00932	0.0623	0.0103	0.0608	0.0103	0.0623	0.00789	0.0745	
8	0.0625	0.102	0.00347	0.0478	0.00742	0.0446	0.00748	0.0444	0.00748	0.0444	0.00800	0.0430	0.00739	0.0450	0.00512	0.0571

Inductances in μH, capacitances in μF for Z_0 = 1 ohm and τ = 1 μs.
For other values, multiply inductance values by Z x τ 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,

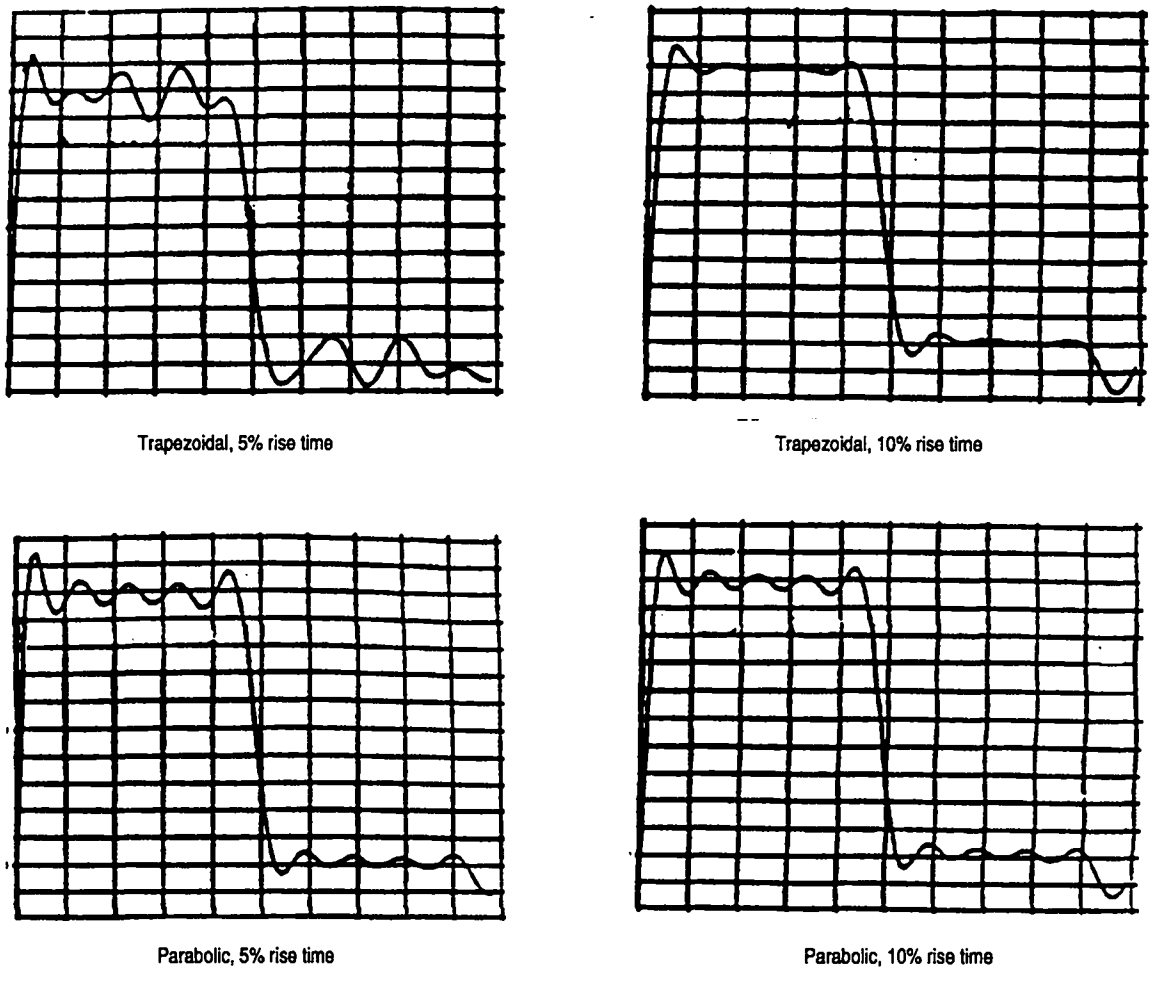

Fig. 9-20. Computer-derived wave shapes for five-element Type C networks (fundamental and third, fifth, seventh and ninth harmonics).

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-

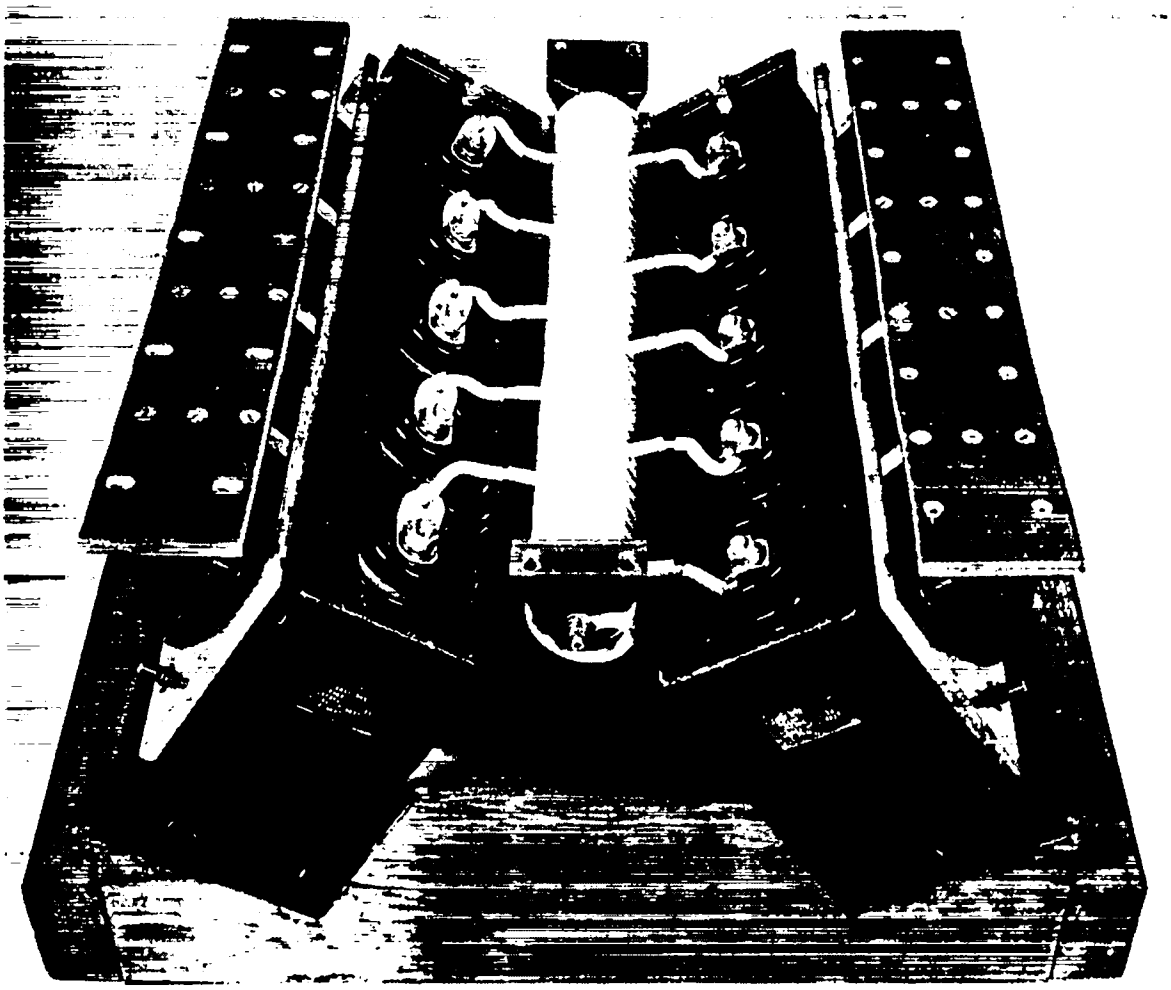

Figure 9-21. A high-performance 10-section Type E pulse-forming network.

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-μs-duration pulses of 13 MW peak power with 0.25-μs 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 μF.

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^2dt, 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

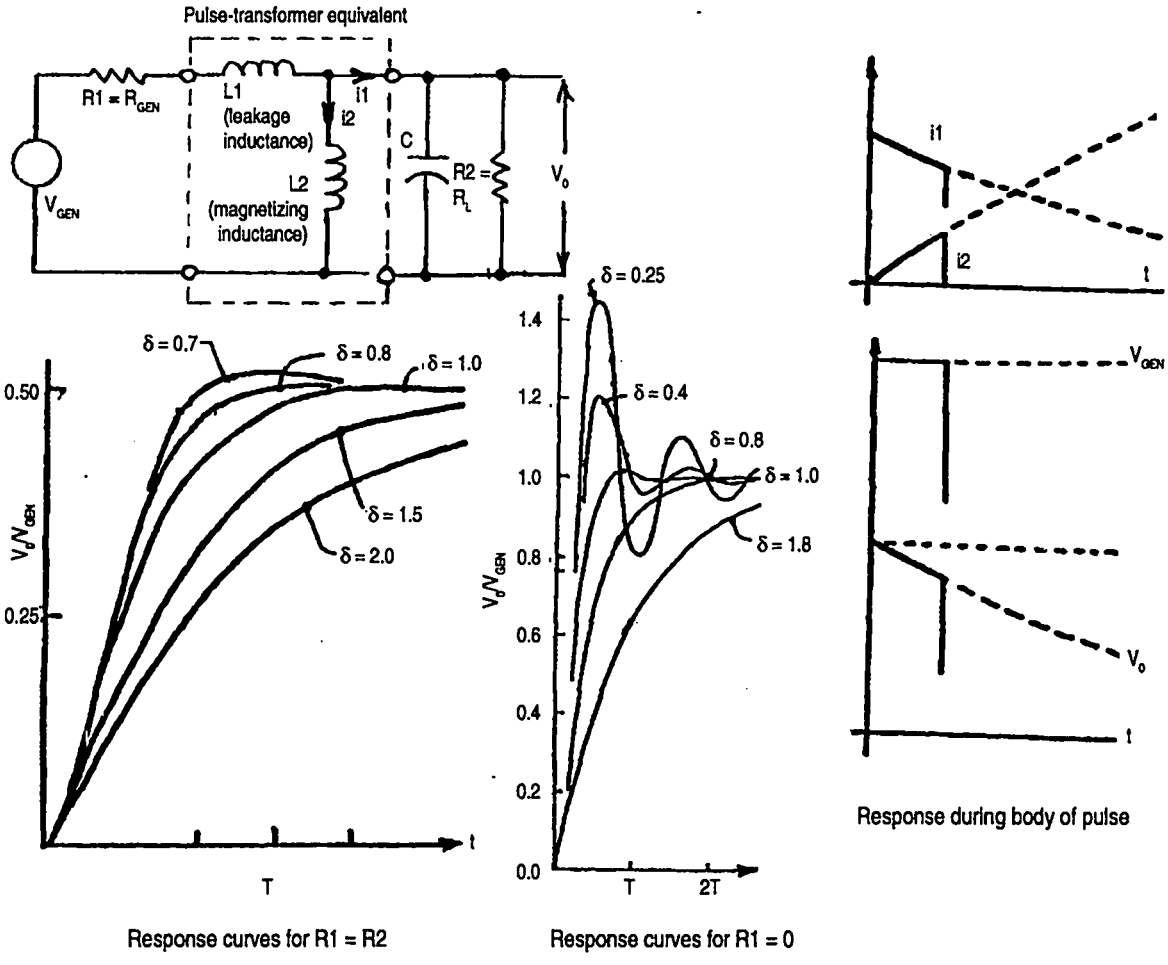

Figure 9-22. Pulse transformer characteristics and how they affect pulse wave shapes.

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

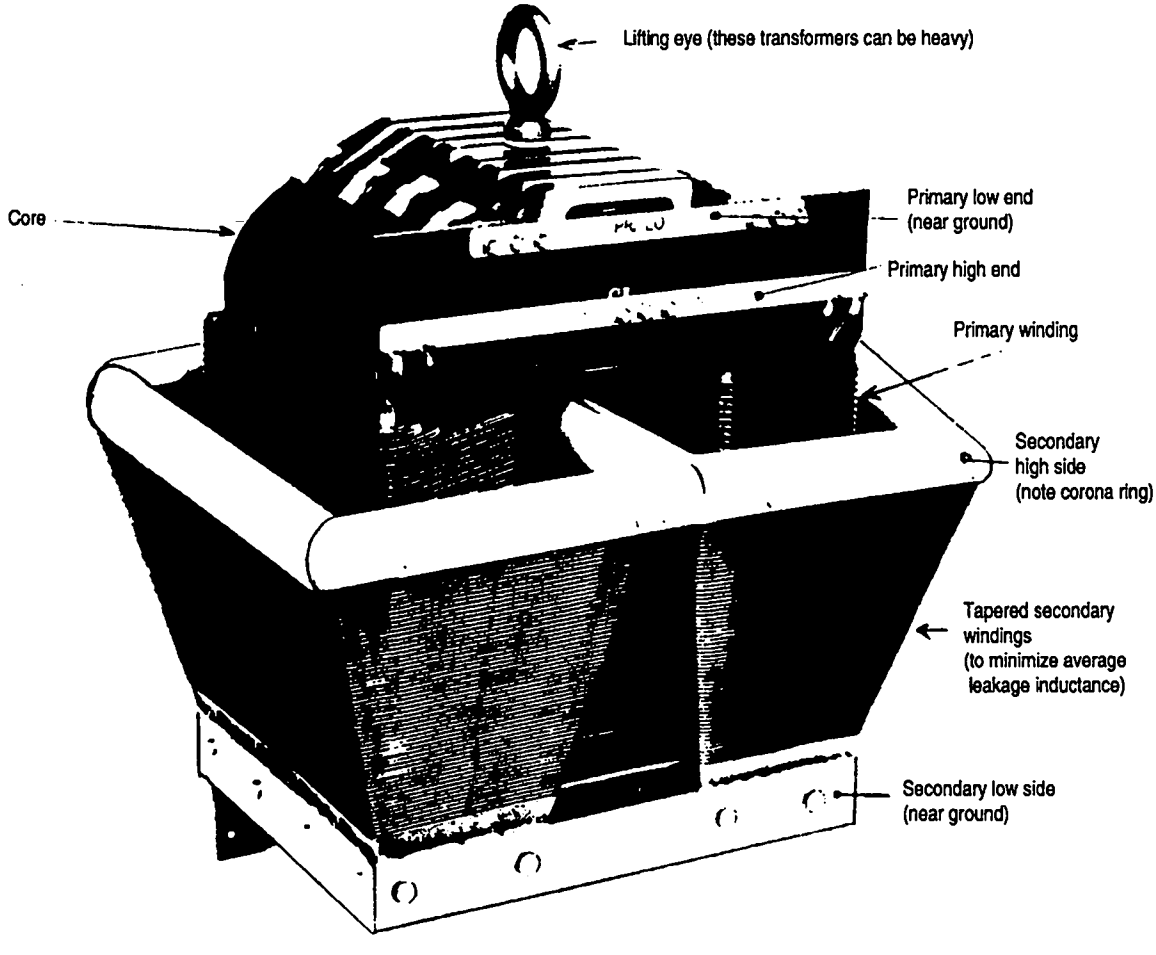

Figure 9-23. A typical high-voltage, high-power output-pulse transformer.

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

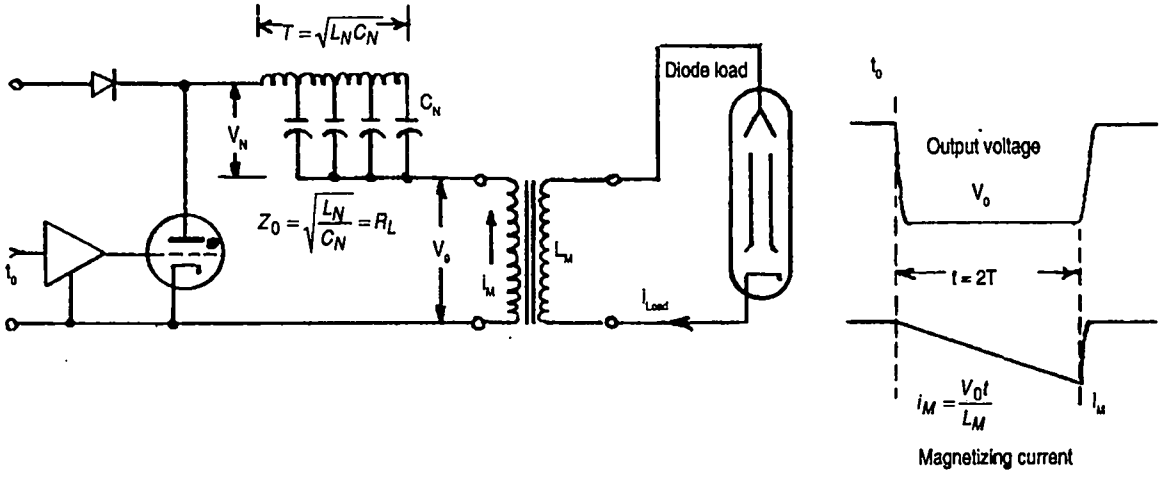

Figure 9-24. What happens to energy stored in pulse-transformer magnetizing inductance.

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

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

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}.

$$W_{NR} = W_M$$

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

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

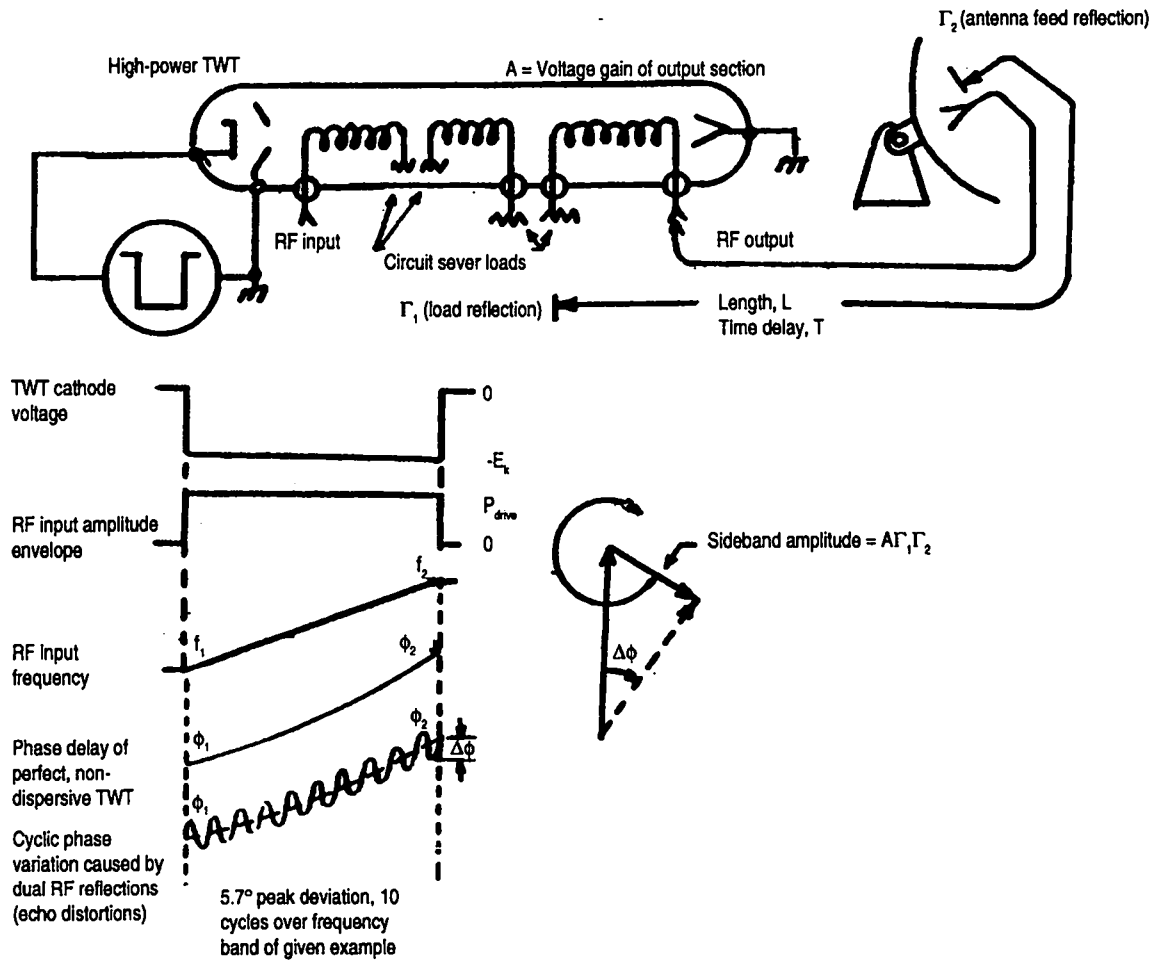

Figure 9-25. Cyclic phase deviation as a function of frequency due to phase-pulling of TWT.

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 x 0.1 x 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 x 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 $f = f_1 + 500\text{x}10^6 t/T$, 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 $500\text{x}10^6\ t^2/2T$.) 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 x 0.1 x 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 x 0.1 x 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 x 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

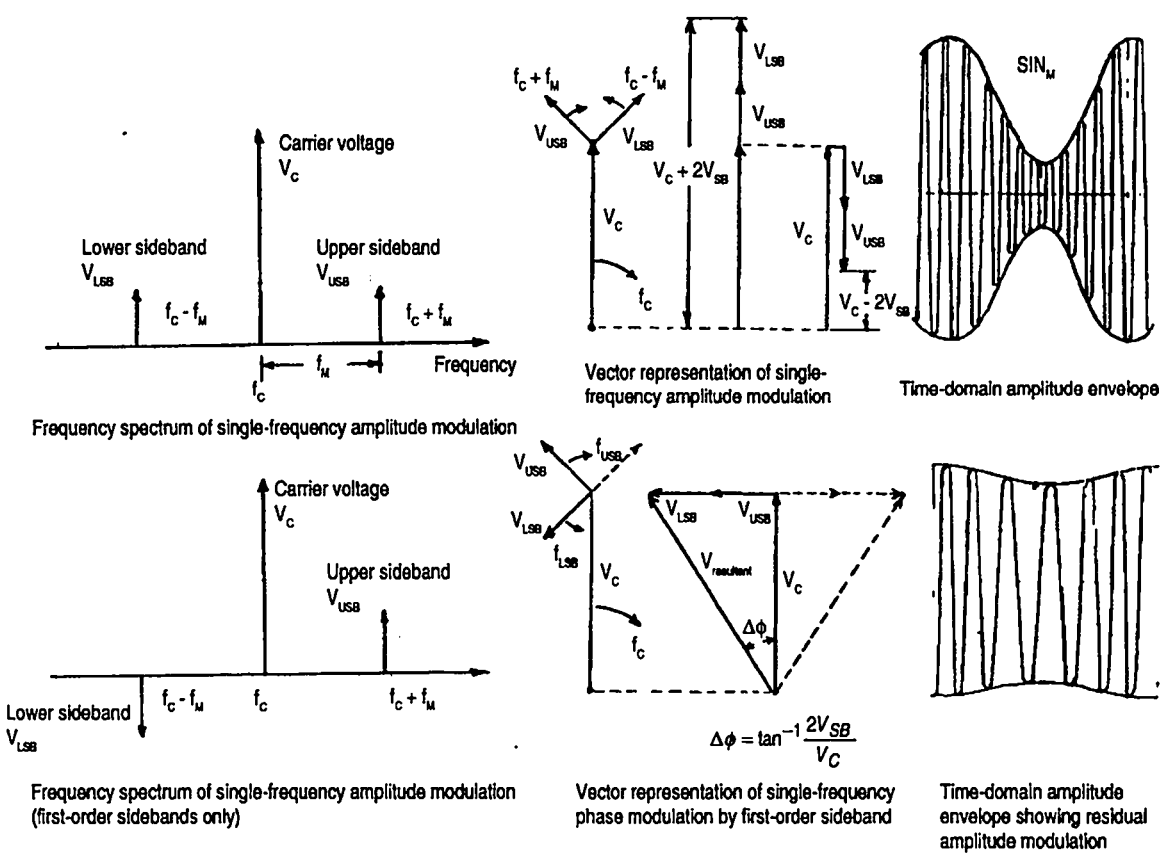

Figure 9-26. Amplitude and phase modulation of an RF carrier by a single modulation frequency.

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

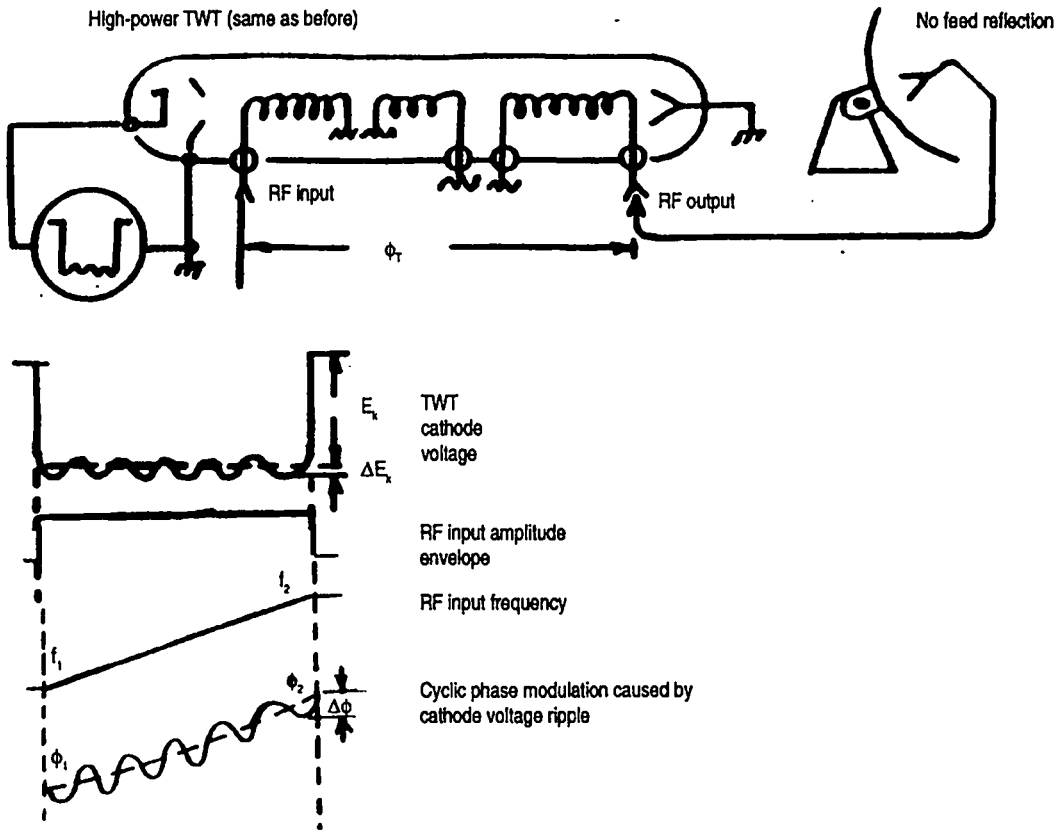

Figure 9-27. Cyclic phase modulation of a TWT through phase pushing.

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 x 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 x 360° (10,800°). Therefore, a 1% change in beam voltage can be expected to produce a phase change of 10,800° x 0.01 x 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	$\frac{\Delta I}{I_0} = 10\frac{\Delta V}{V_0}$	$\frac{\Delta n}{n_0} = 0\frac{\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta n}{n_0}$ $= 11\frac{\Delta V}{V_0}$	$\frac{\Delta\phi}{\phi_0} = 0.4°/1\%\,\Delta I / I_0$ $\Delta\phi = 4°/1\%\,\Delta V / V_0$
Klystron Amplifier (Linear Beam)	$\frac{\Delta I}{I_0} = 1.5\frac{\Delta V}{V_0}$	$\frac{\Delta n}{n_0} = 1.5\frac{\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta n}{n_0}$ $= 4\frac{\Delta V}{V_0}$	$\frac{\Delta\phi}{\phi_0} = \frac{1}{2}\,\Delta V / V_0$ $\Delta\phi = 10°\text{–}15°/1\%\,\Delta V / V_0$ (See note 1)
TWT Amplifier (Linear Beam)	$\frac{\Delta I}{I_0} = 1.5\frac{\Delta V}{V_0}$	$-\frac{0.5\Delta V}{\Delta V_0} < \frac{\Delta n}{n_0} < +\frac{0.5\Delta V}{V_0}$	$\frac{\Delta P_{RF}}{P_{RF}} = \frac{\Delta V}{V_0} + \frac{\Delta I}{I_0} + \frac{\Delta n}{n_0}$ $= 2 - 3\frac{\Delta V}{V_0}$	$\frac{\Delta\phi}{\phi_0} = \frac{1}{3}\,\Delta V / V_0$ $\Delta\phi = 10°\text{–}15°/1\%\,\Delta V / V_0$ (See note 2)

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

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%

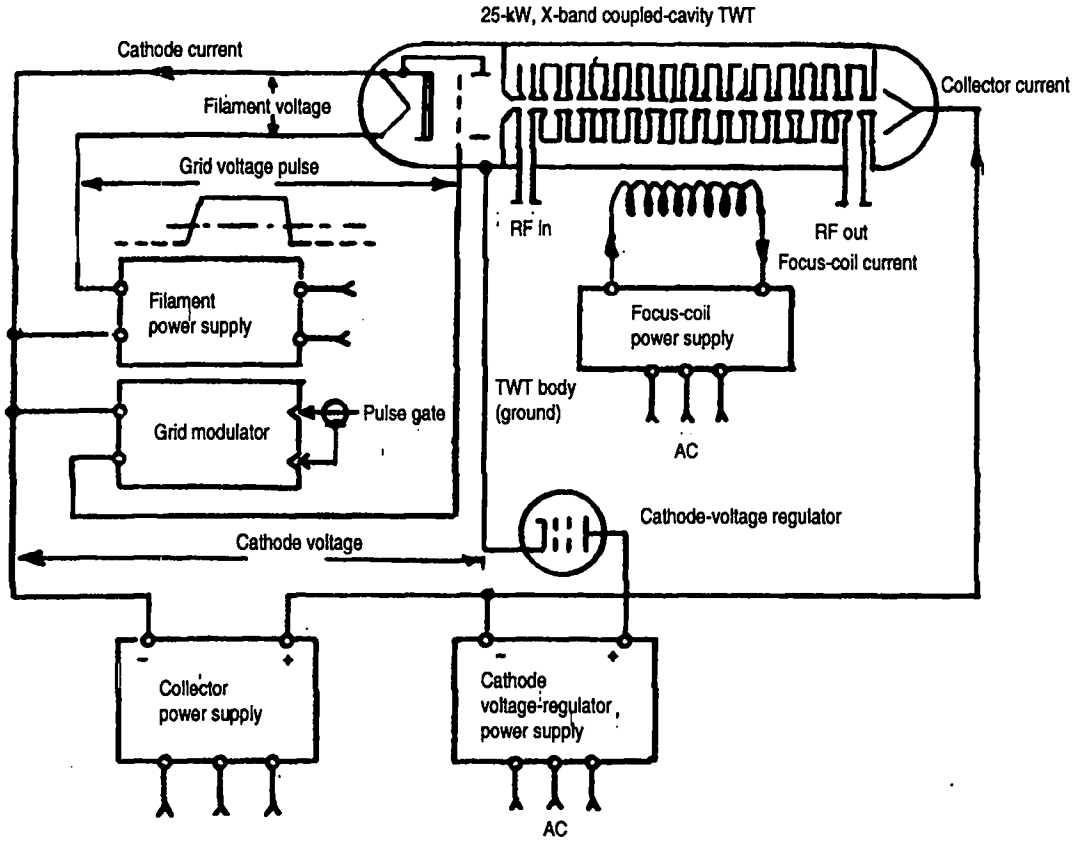

Figure 9-29. TWT electrical inputs that affect amplitude- and phase-pushing (see Table 9-1).

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

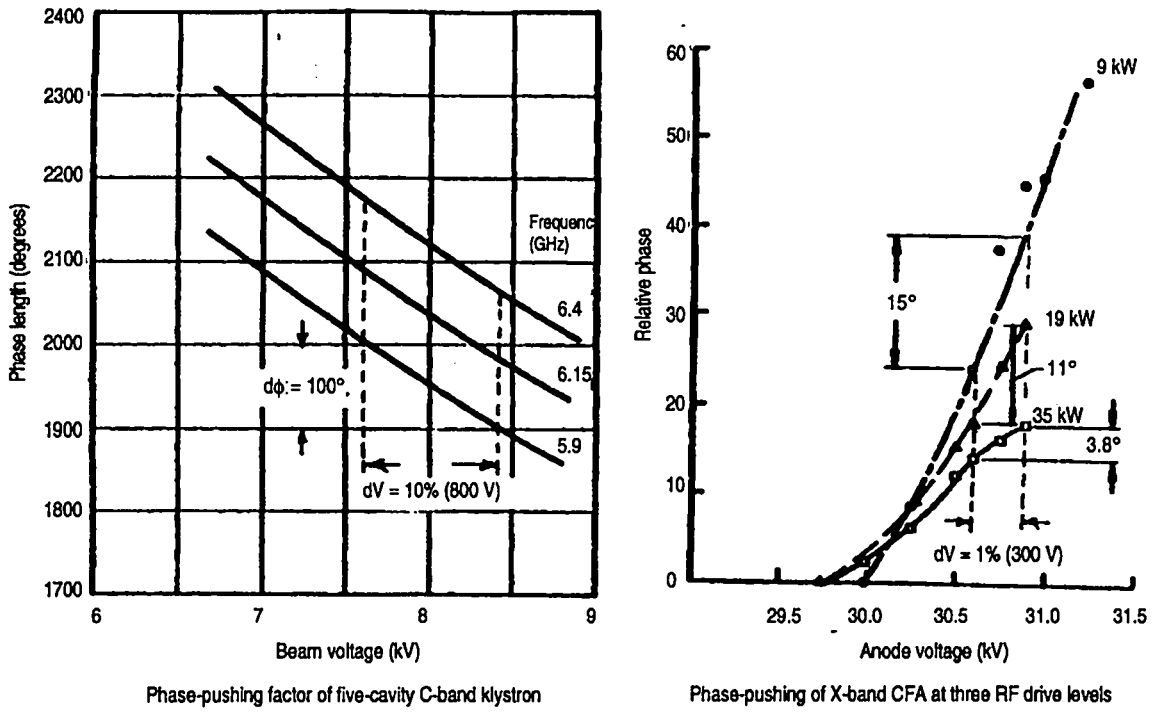

Figure 9-30. Measured phase-pushing factors for klystron and crossed-field amplifiers.

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

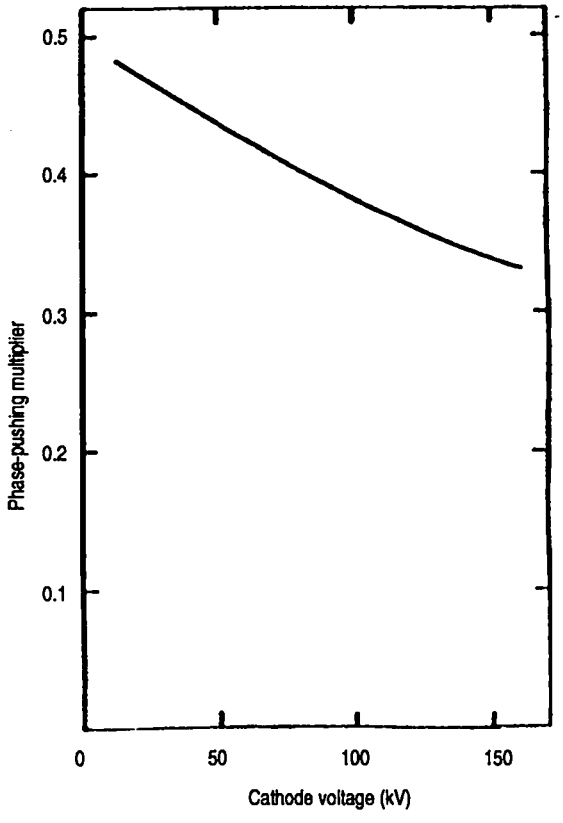

Figure 9-31. Phase-pushing multiplier plotted against cathode voltage.

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

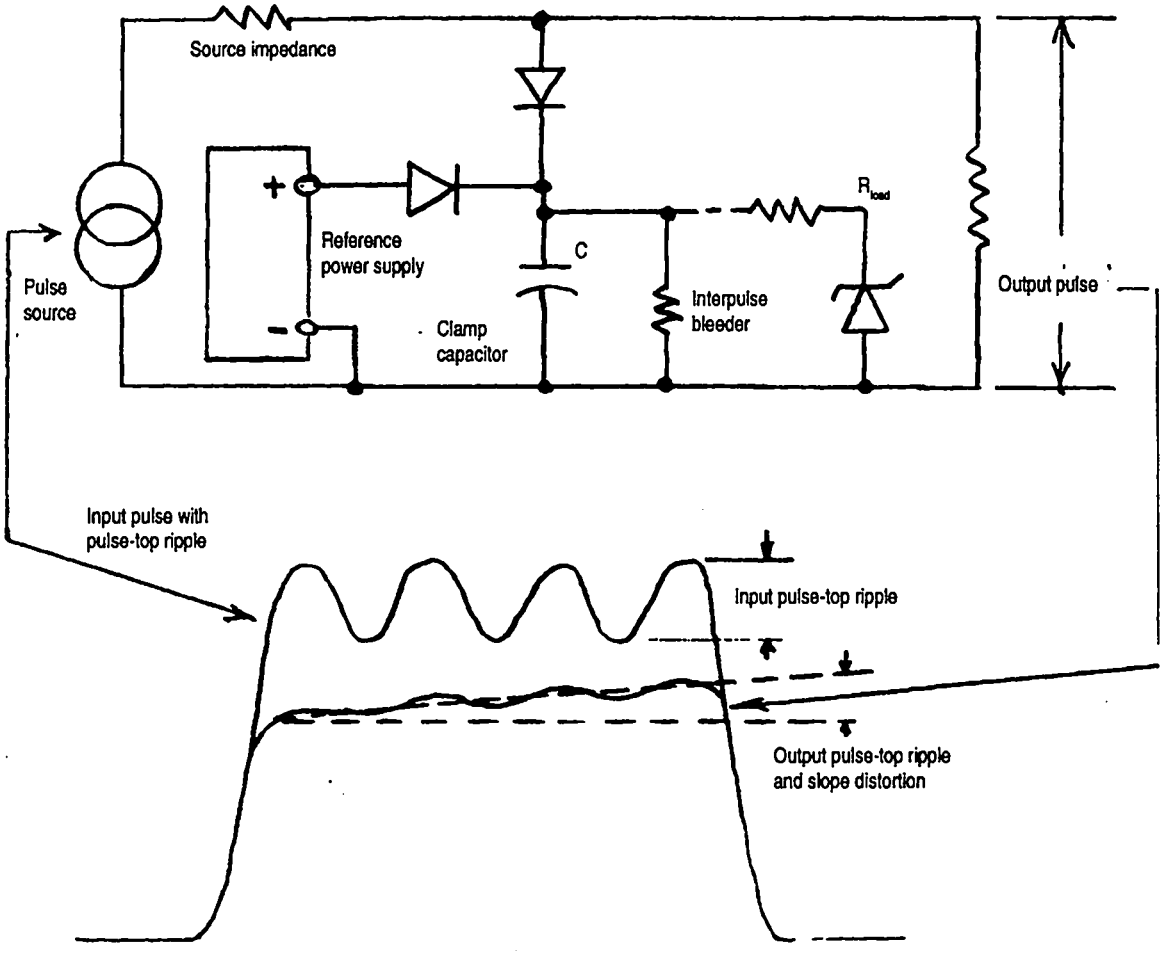

Figure 9-32. Attenuation of pulse-top ripple with a clamping circuit.

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

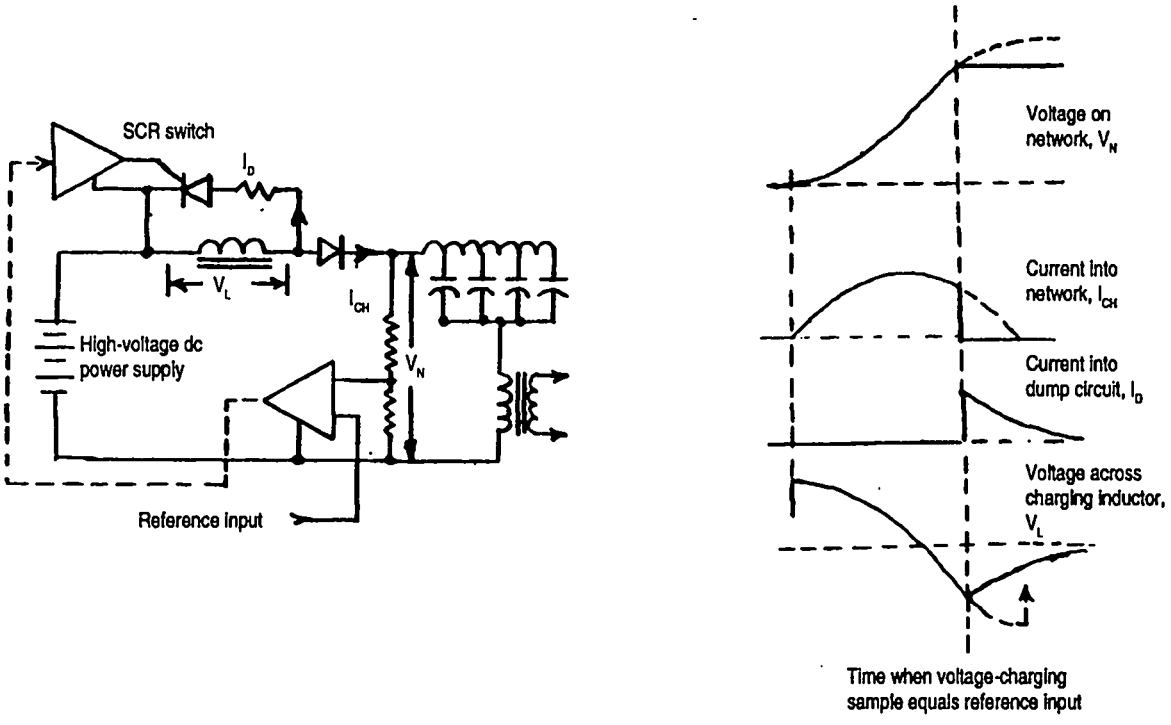

Figure 9-33. Pulse-to-pulse network charging regulation by deQing.

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 thyratron

The hydrogen thyratron 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 "thyratron" and "hydrogen thyratron" are often

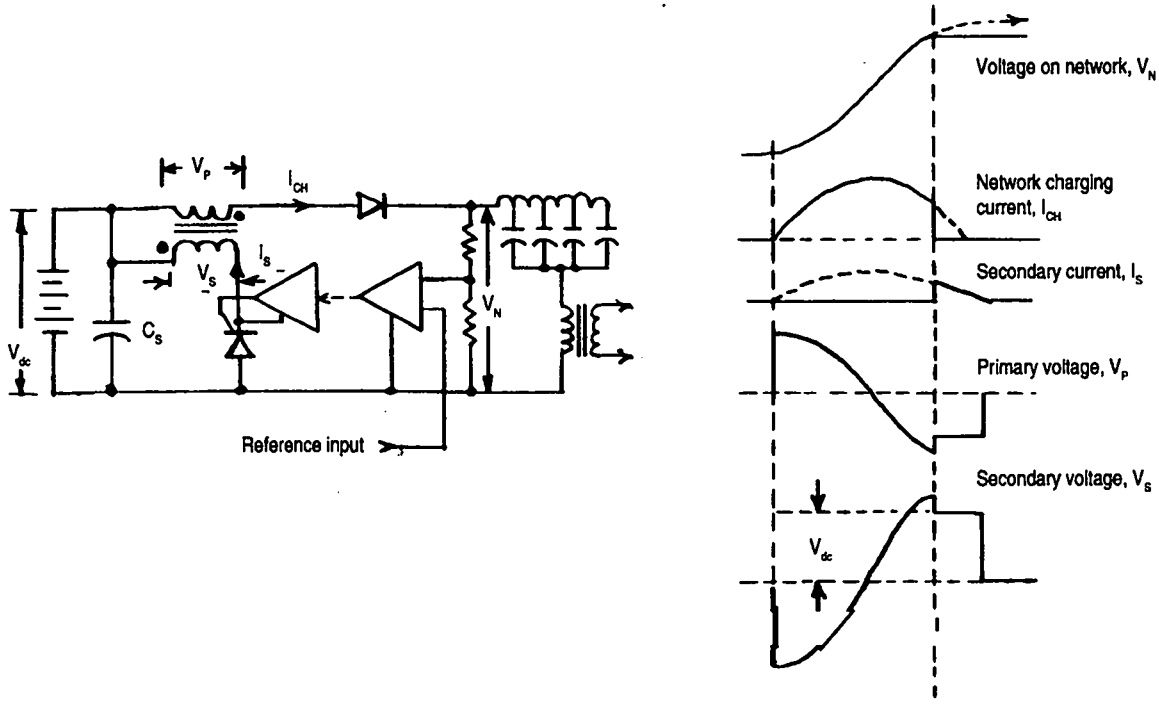

Figure 9-34. High-efficiency pulse-to-pulse regulation through charge recycling.

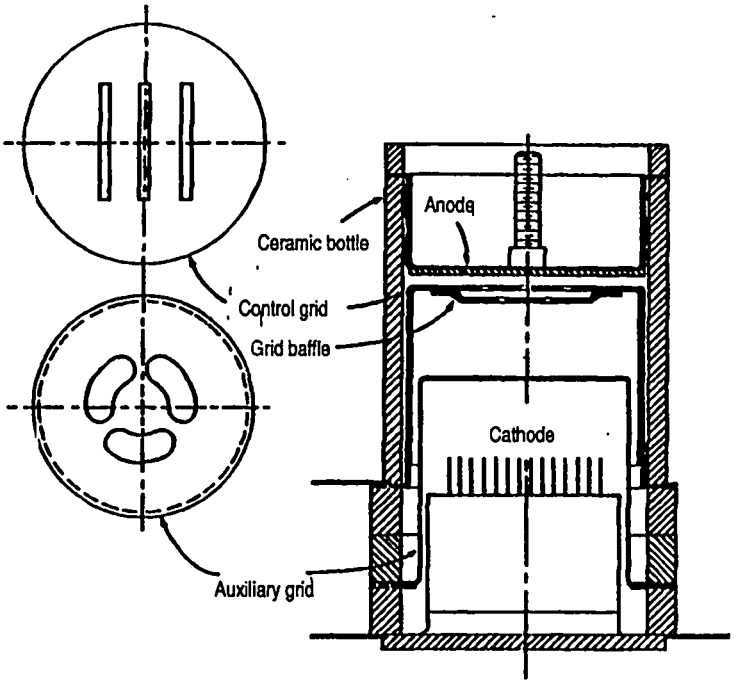

Figure 9-35. Cross-section view of hydrogen thyratron.

used synonymously, thyratron is the more generic term, referring to the early devices that were filled with mercury, argon, and other gases. The charge carriers in a thyratron 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 thyratrons 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, thyratron. 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 thyratron. This is because hydrogen thyratrons are designed to be high-power switches. They conduct hundreds, or even thousands, of amperes and, in the preconduction 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-

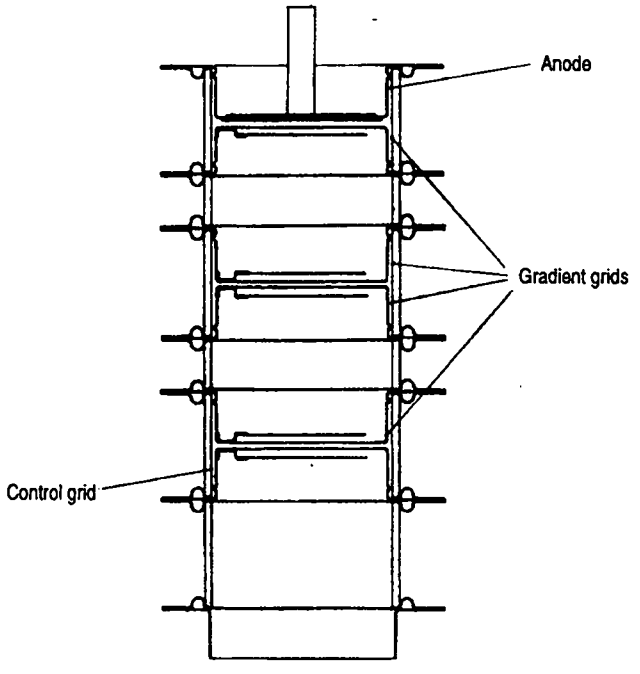

Figure 9-36. Increasing thyratron voltage hold-off by means of voltage-gradient grids.

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 thyratron, 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 thyratron 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 thyratrons 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-

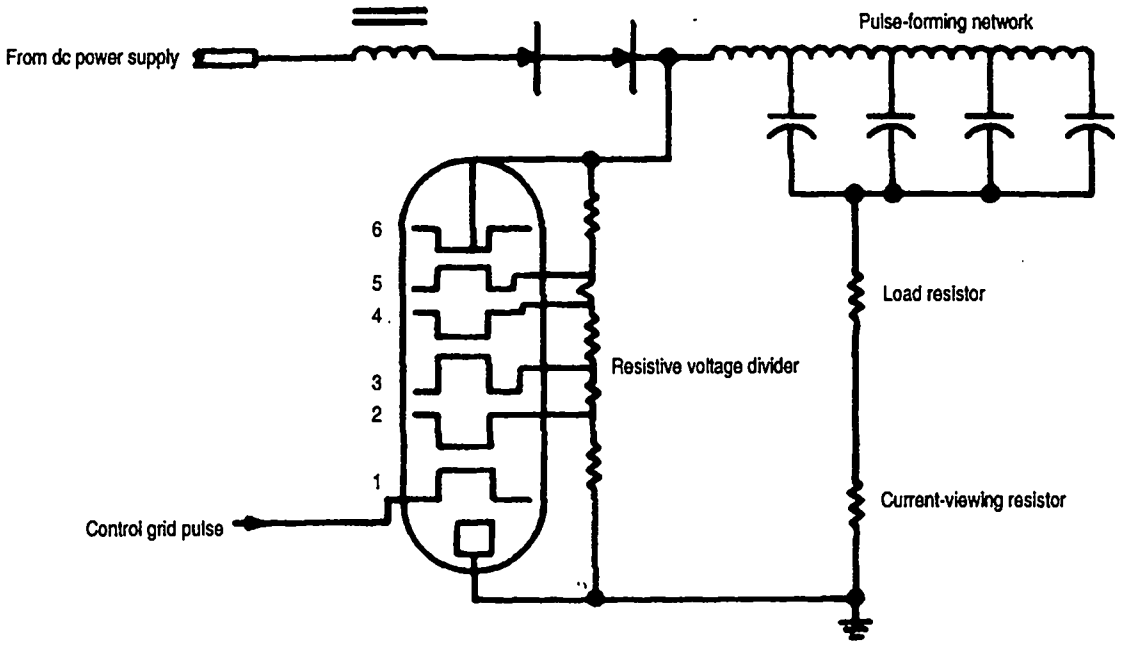

Figure 9-37. Circuit arrangement using gradient-grid thyratron.

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 thyratron 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, H_2. 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 μs 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 μs 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.

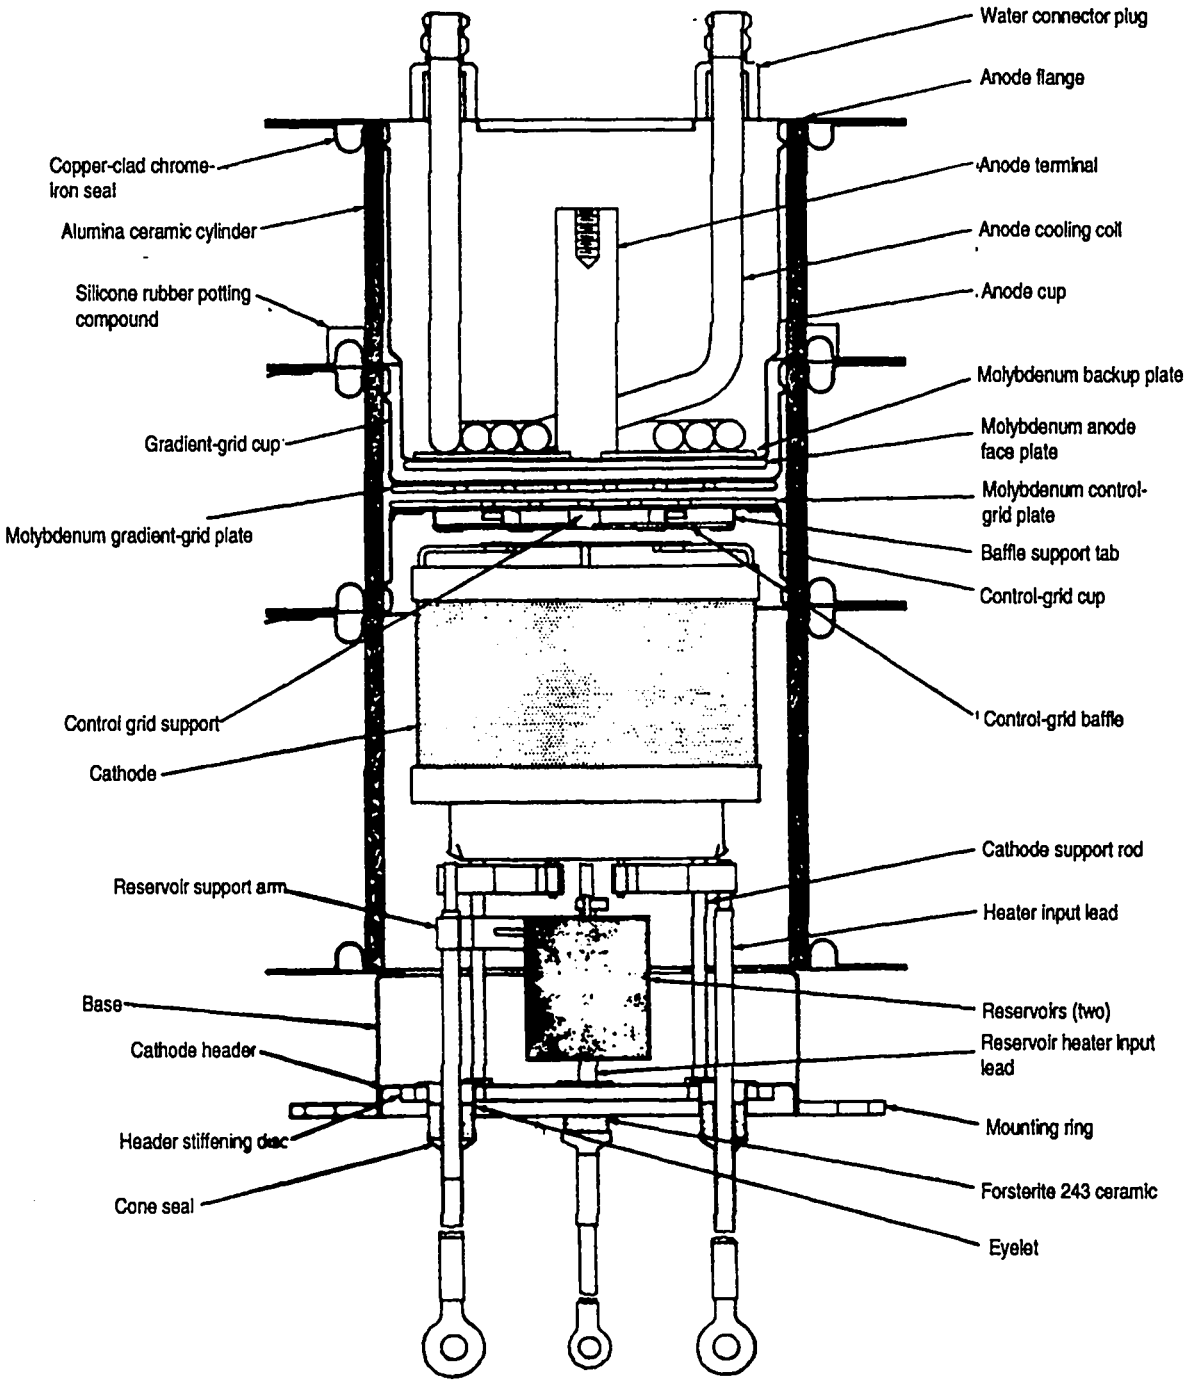

Figure 9-38. Internal construction of a high-power hydrogen thyratron.

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 thyratron 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 thyratron—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 thyratron had a voltage-enhancing gradient grid, two hydrogen reservoirs, a 200-cm^2 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 $55x10^9$. (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 $52x10^9$, however. This was the only performance characteristic close to the tube's rating.

Very much the same geometry was used in a thyratron with even higher performance specifications: the 40-kV, 1-MW-average-power MAPS-40 thyratron. 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 $400x10^9$. 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 cm^2.

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-

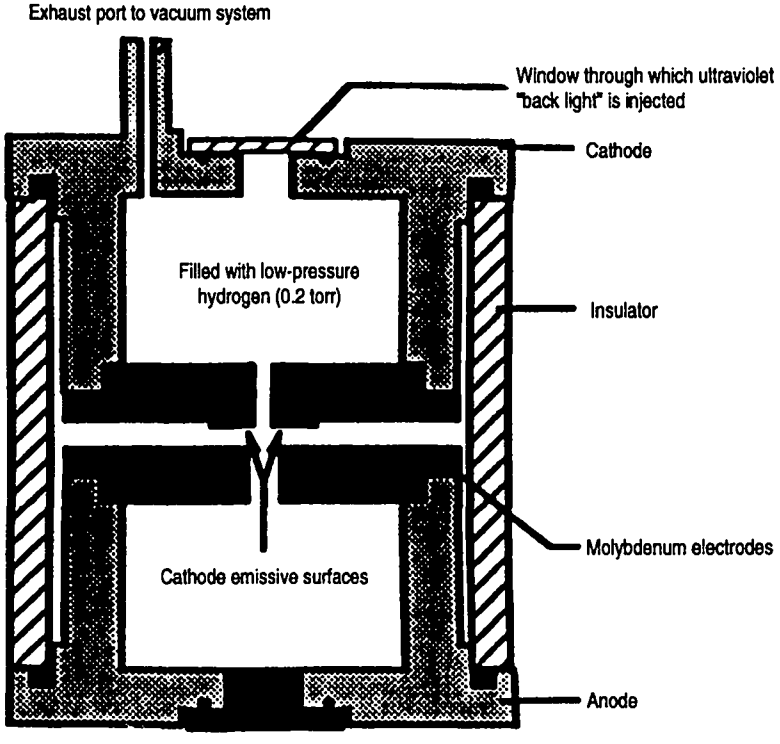

Figure 9-39. Cross-section view of "back-lit" thyratron.

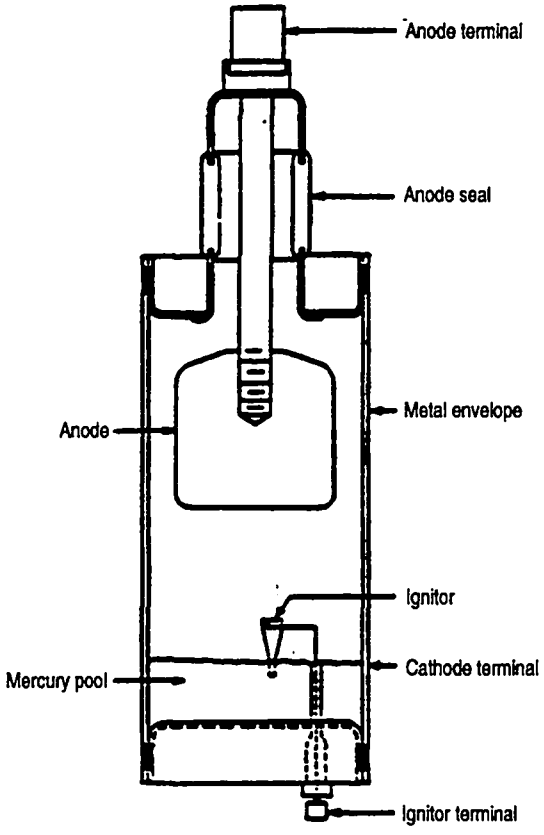

Figure 9-40. Cross-section view of mercury-vapor ignitron.

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" thyratron. Filled with low-pressure hydrogen, it, like the conventional thyratron, 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.

Thyratrons 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 thyratron is the cold-cathode, instant-start thyratron. 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 thyratron. 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 thyratron 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 thyratron 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-μs 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.

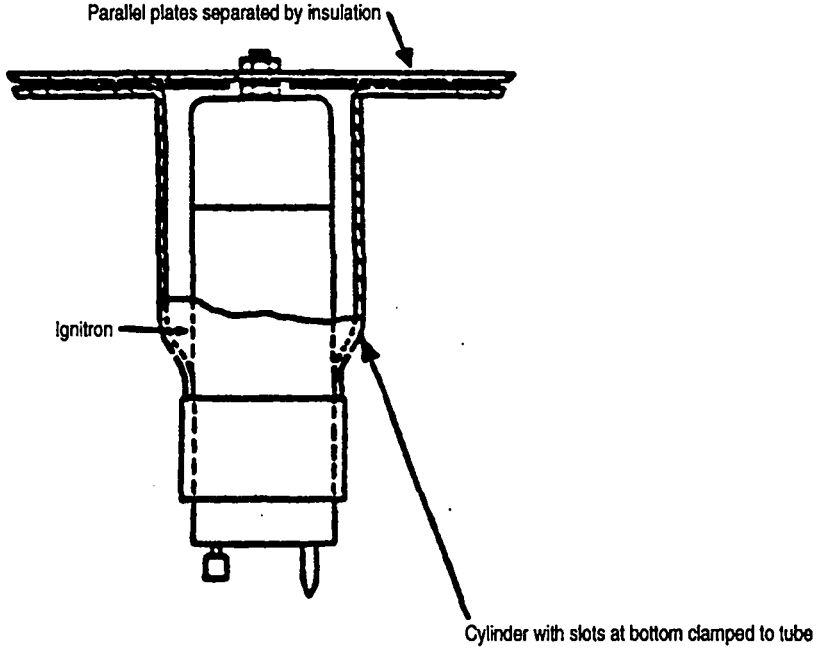

Figure 9-41. Coaxial mounting for ignitron switch.

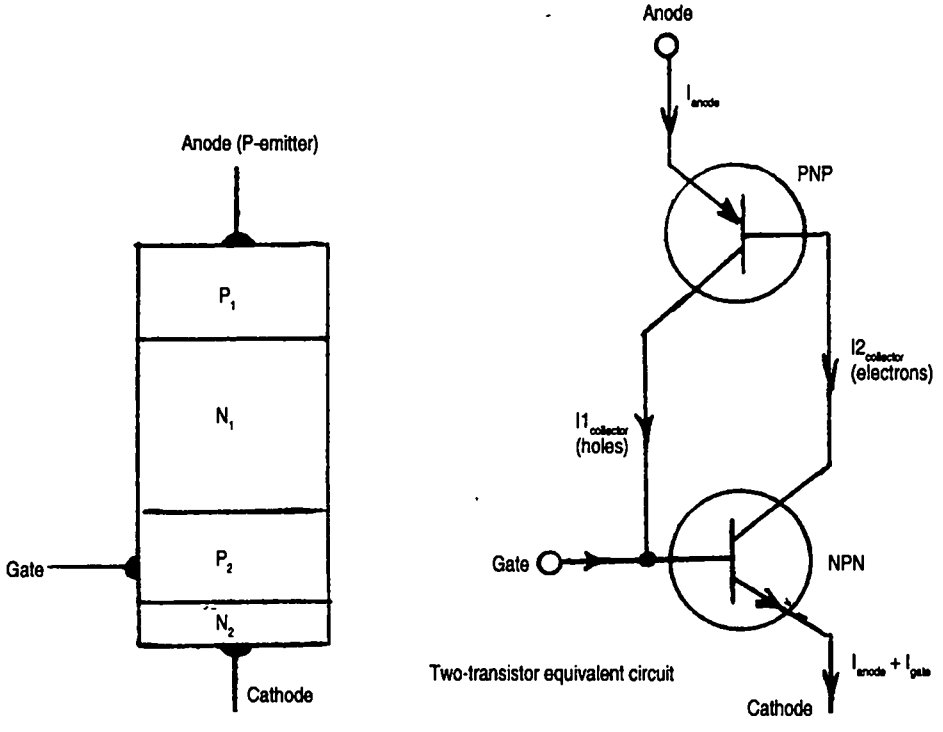

Figure 9-42. The thyristor, or silicon controlled rectifier (SCR).

9.4.3 The thyristor

The thyristor (a term derived from the words "thyratron" 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 thyratron. 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 thyratron, 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 thyratron. 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 μs, and *dv/dt* , typically 100 V/μs 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 thyratron. (A 50-kV thyratron, 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 thyratron 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 overvoltaging 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/μs rate-of-rise is considered the minimum acceptable. This configuration gives peak discharge currents greater than 500 A at *di/dt* of 2000 A/μs and at pulse durations over 50 μs. For the larger 0.9-in.-diameter fusion, 15000 V/μs is minimum for peak discharge currents up to 5000 A, 2500 A/μs rate-of-rise, and at pulse durations up to 20 μs. 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 thyratron, 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 thyratron, 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 thyratron

Although it is presumptuous, considering the rich history of the thyratron, 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 thyratron switch capable of handling the full power required. However, a single ignitron switch could, and it was given serious consideration. Nevertheless, the thyratron prevailed as the switch of choice,

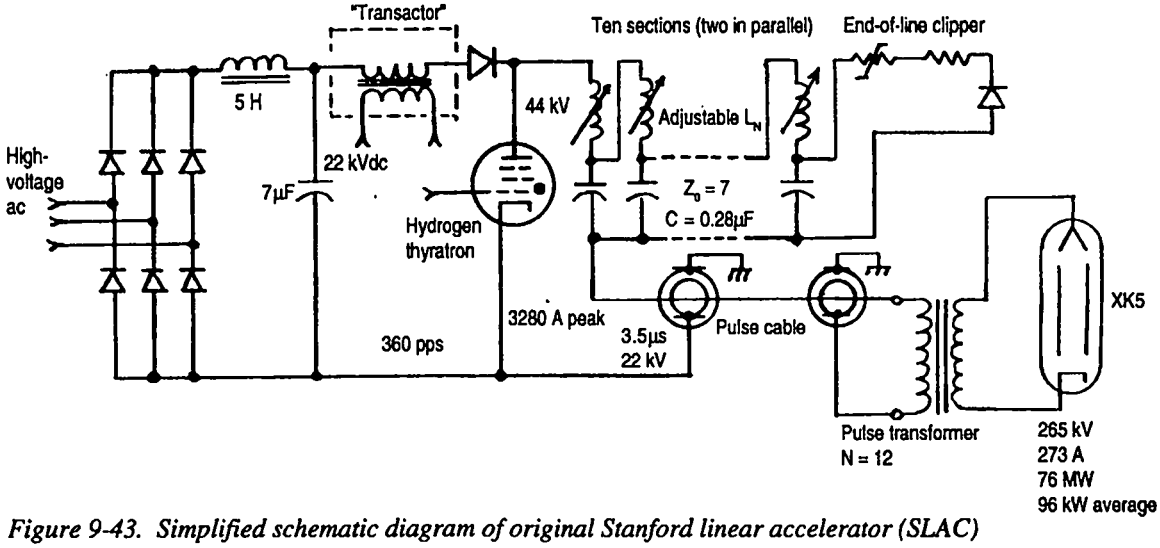

Figure 9-43. Simplified schematic diagram of original Stanford linear accelerator (SLAC) pulse modulator.

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-μs 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 μF. Reverse thyratron 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 nonlinear 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 thyratron 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 thyratron 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

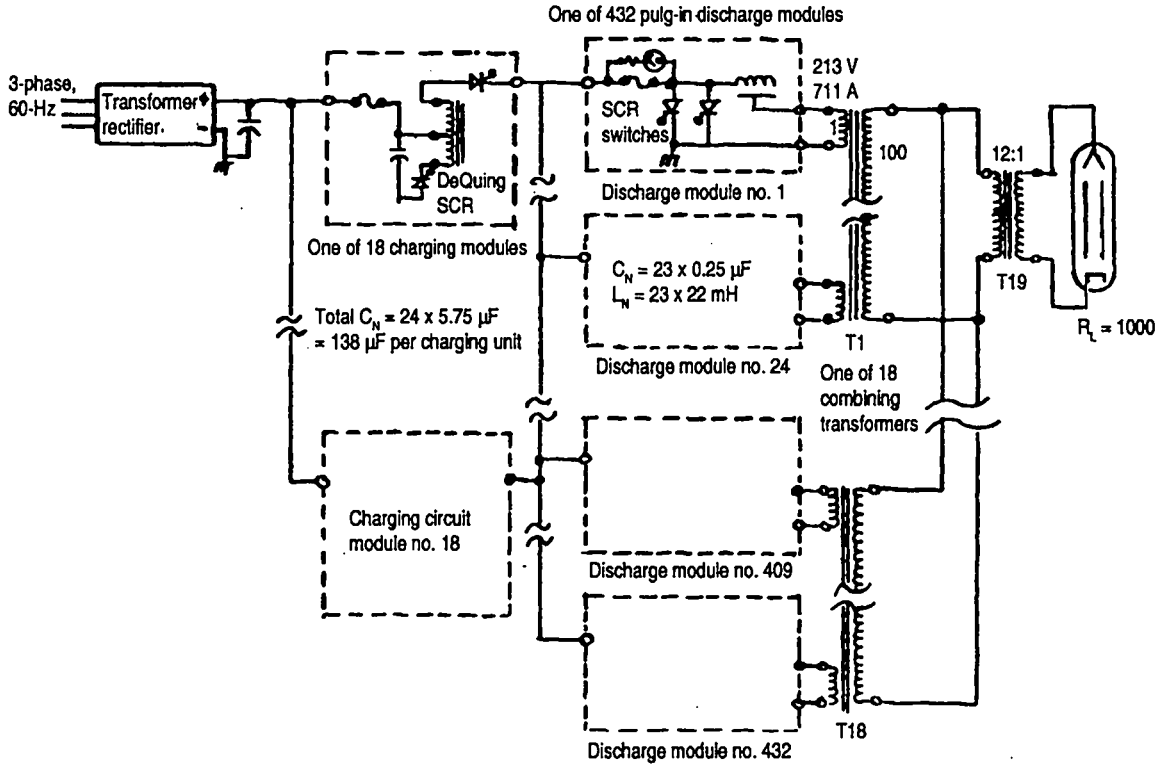

Figure 9-44. Solid-state version of original SLAC pulse modulator.

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 thyratron 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 deQing 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

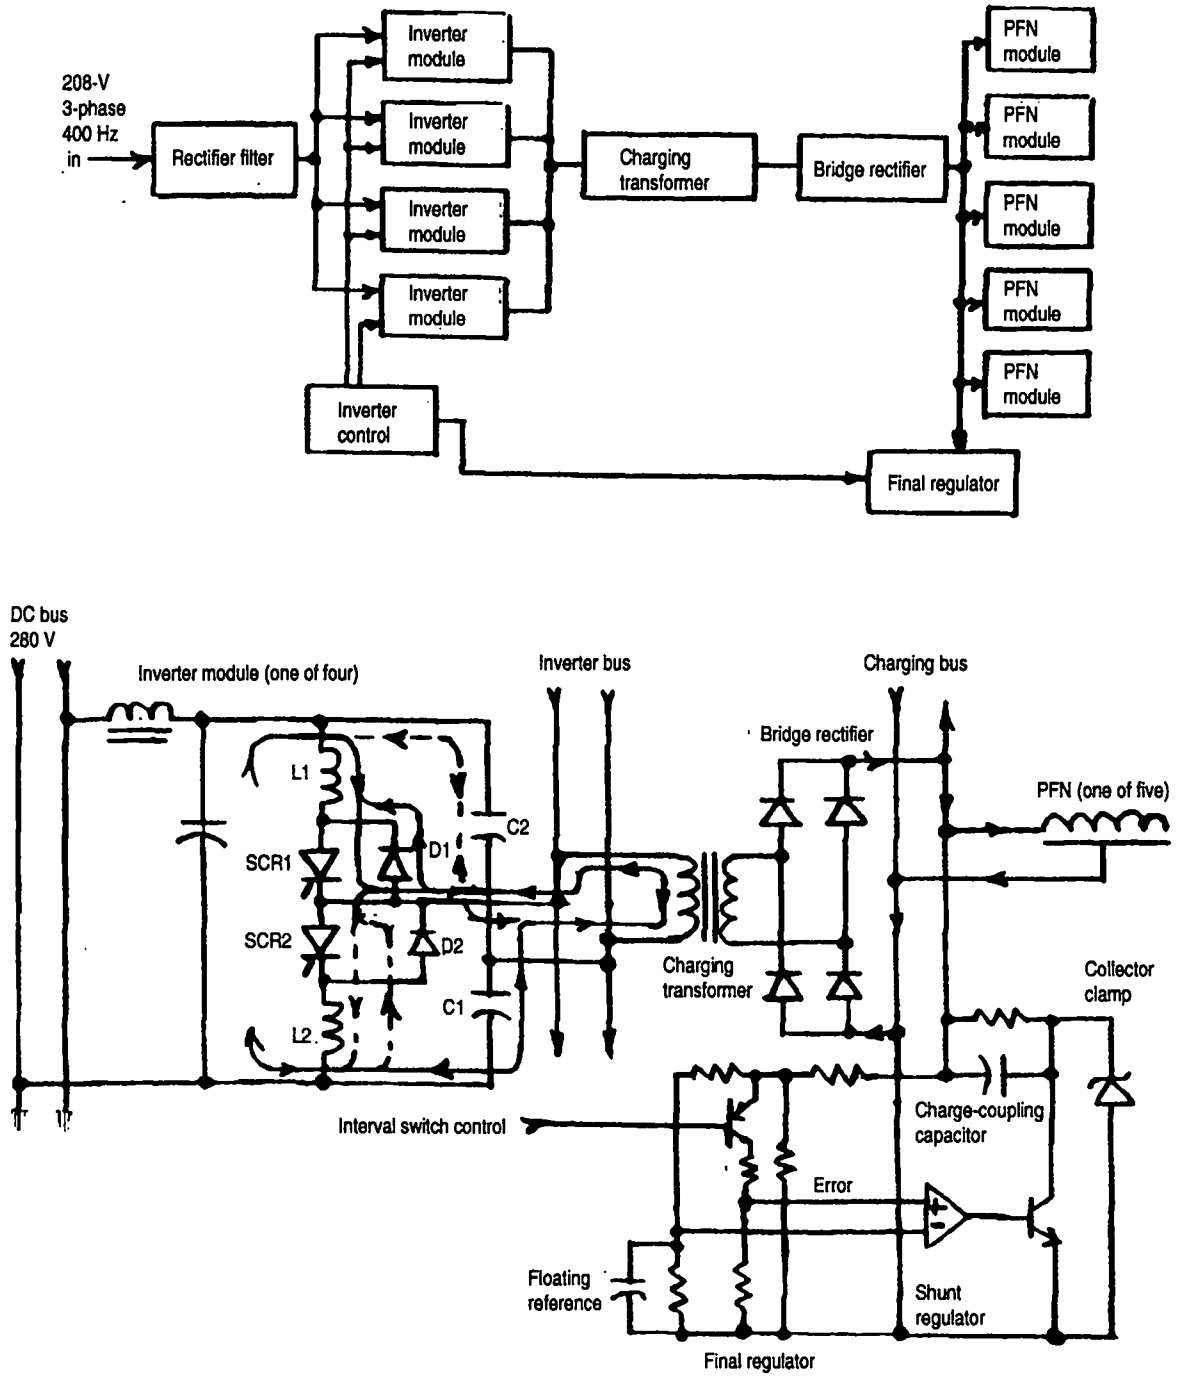

Figure 9-45. Charging circuit of high-performance solid-state line pulser.

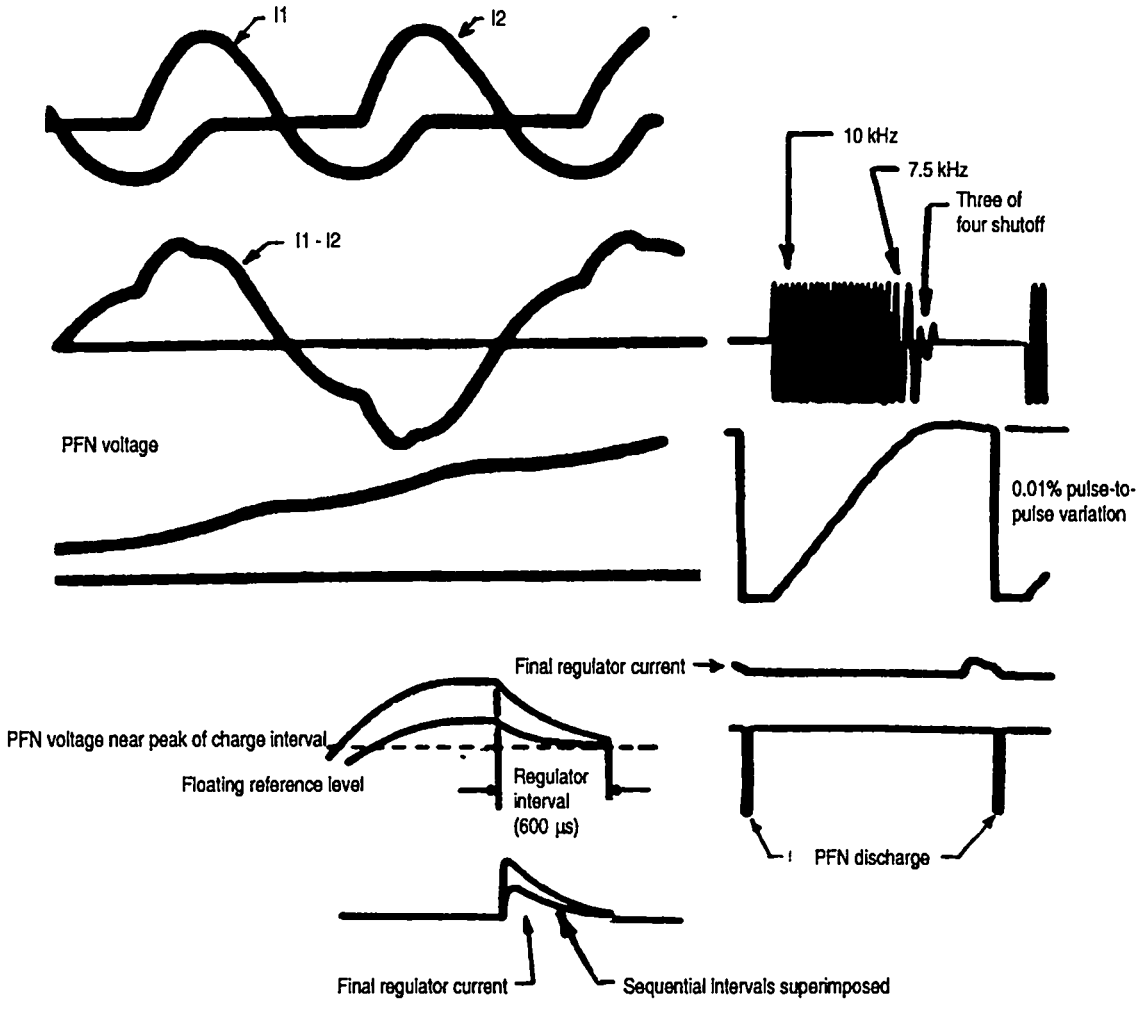

Figure 9-46. Solid-state modulator charging-circuit waveforms.

a thyratron. 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 μs 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.

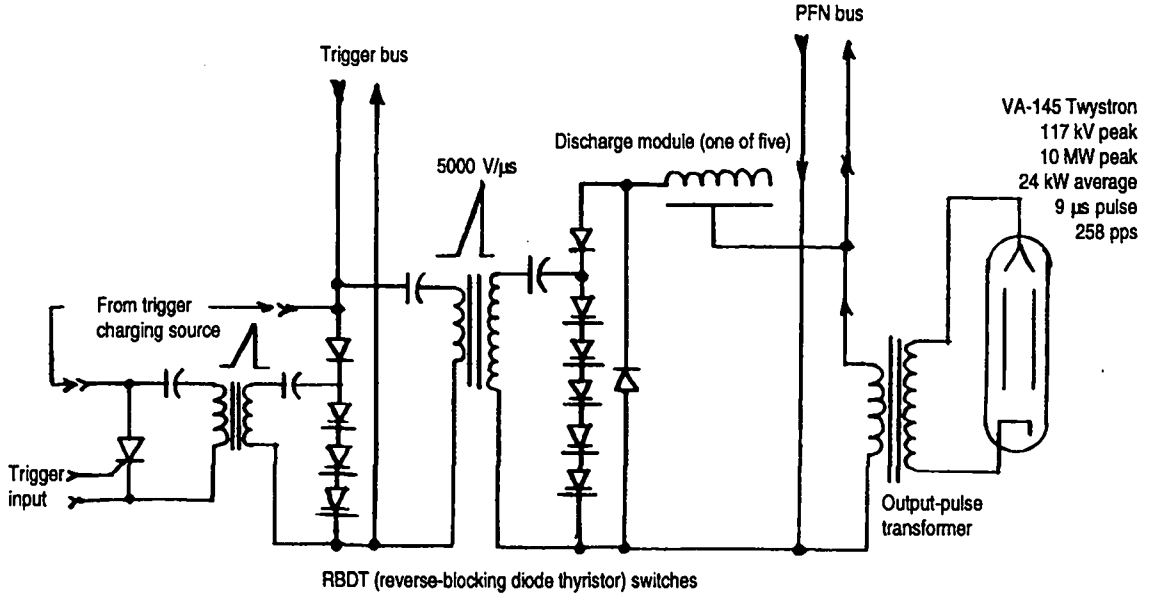

Figure 9-47. Solid-state modulator discharge circuits, using RBDTs as output switches.

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-μs 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 μs. (Needless to say, a PFN having a delay time of 1000 μs 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 *CR2* and the power-supply storage fused capacitor bank, *C1*, 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, *CR3* and *R1*.

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.

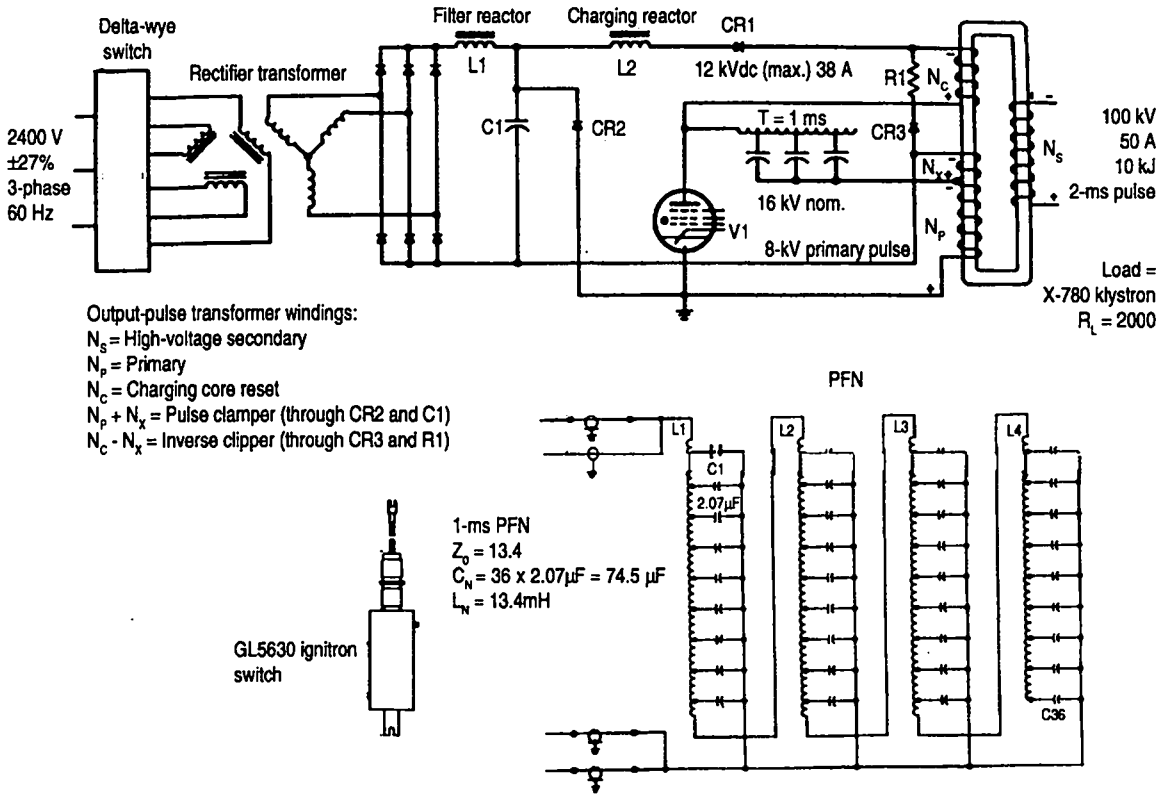

Figure 9-48. Long-pulse line-type modulator using ignitron output switch.

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-μs 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

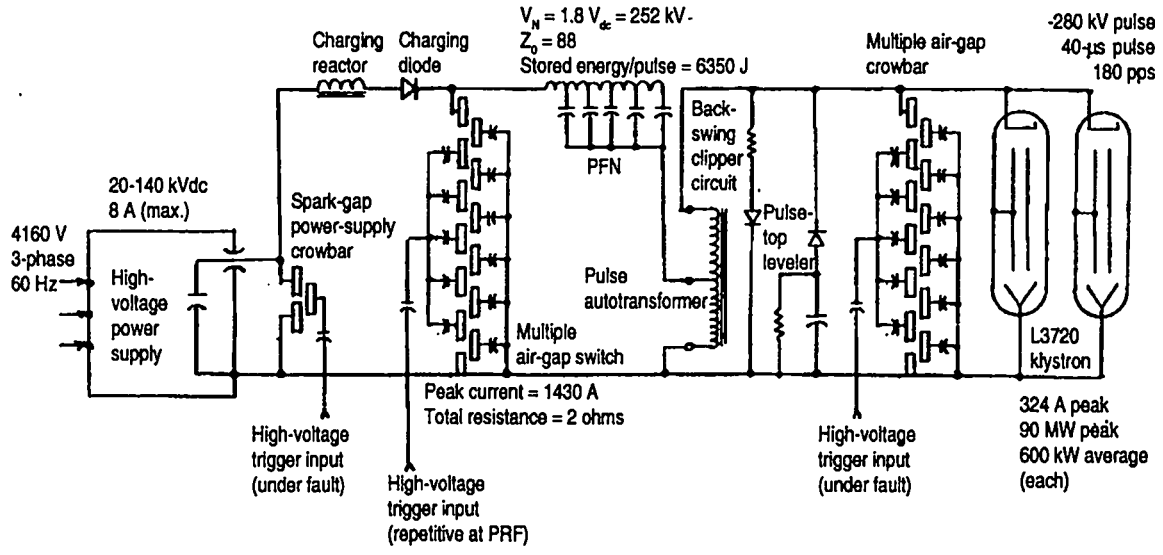

Figure 9-49. Super-power line-type modulator using triggered air-gap output switch.

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

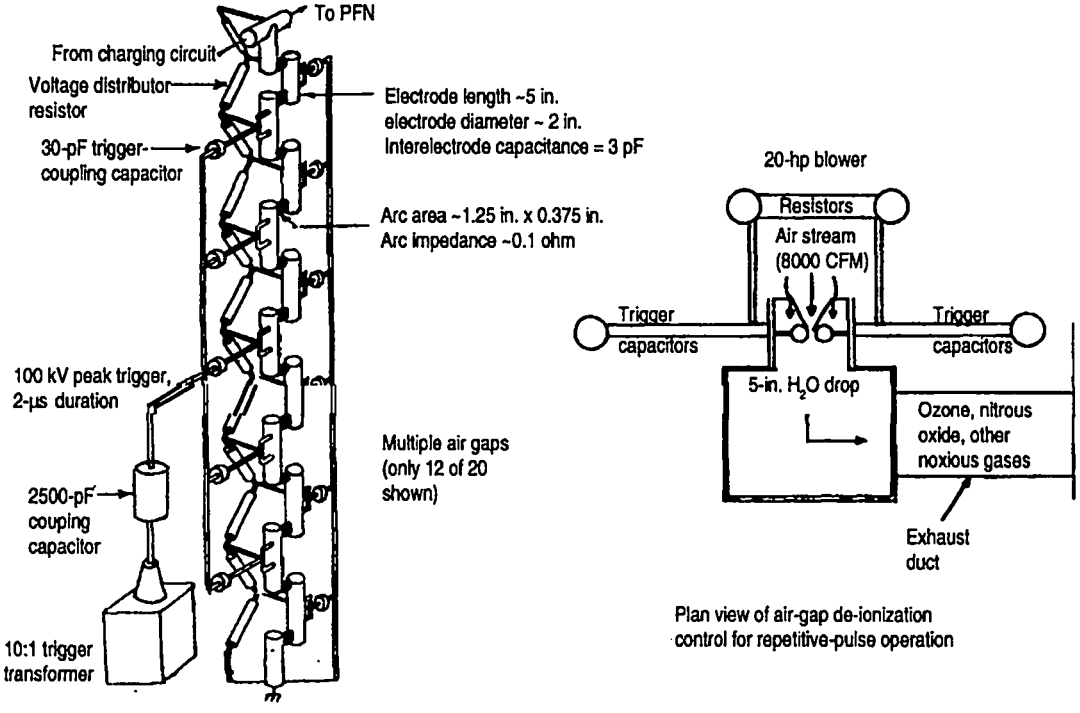

Figure 9-50. Details of multiple-stage triggered air-gap output switch.

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 ft^3/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.