

Switching Regulators for Poets

A Gentle Guide for the Trepidatious

Jim Williams

The above title is not happenstance and was arrived at after considerable deliberation. As a linear IC manufacturer, it is our goal to encourage users to design and build switching regulators. A problem is that while everyone agrees that *working* switching regulators are a good thing, everyone also agrees that they are difficult to get working. Switching regulators, with their high efficiency and small size, are increasingly desirable as overall package sizes shrink. Unfortunately, switching regulators are also one of the most difficult linear circuits to design. Mysterious modes, sudden, seemingly inexplicable failures, peculiar regulation characteristics and just plain explosions are common occurrences. Diodes conduct the wrong way. Things get hot that shouldn't. Capacitors act like resistors, fuses don't blow and transistors do. The output is at ground, and the ground terminal shows volts of noise.

Added to this poisonous brew is the regulator's feedback loop, sampled in nature and replete with uncertain phase shifts. Everything, of course, varies with line and load conditions—and the time of day, or so it seems. In the face of such menace, what are Everyman and the poets to do?

The classic approach is to seek wisdom. Substantial expertise exists but is concentrated in a small number of corporate and academic areas. These resources are not readily accessed by Everyman and some cynics might suggest that they are deliberately protected by a self-serving priesthood. A glance through conference proceedings and published literature yields either a storm of mathematics or absurdly coy and simple little block diagrams that make everything look just so easy. Either way, Everyman loses. And the poets don't even get to try.

Something to think about is that most people who want switching regulators don't need 98.2% efficiency or 100W/cubic inch. They aren't trying to get tenure and don't care about inventing a new type of circuit. What they want are concepts directly applicable to construction of working circuits with readily-available parts. Thus equipped, Everyman can build and sell useful products, presumably buy more components and everyone's interests (not incidentally, including ours) are served.

As author, I must confess that I am more poet than switching regulator designer, and my poetry ain't very good. Before this effort, my enthusiasm level for switchers resided somewhere between trepidation and terror. This position has changed to one of cautiously respectful optimism. Several things aided this transformation and significantly influenced this publication. The "encouragement" of the Captains of this corporation, emphasized over the last year at increasingly insistent levels, constituted one form of inspiration. Conversations with users (or people who wanted to be) provided more valuable perspective and strength in the knowledge that I was not alone in my difficulties with switchers.

At the circuit level, a significant decision was to employ standard, off-the-shelf magnetics exclusively.¹ This policy was driven by the observation that the majority of problems encountered with switchers centered around inductive components. This approach almost certainly prevents precisely-optimized performance and may horrify some veteran switcher designers. It also eliminates inductor construction uncertainties, saves time and greatly increases the likelihood of getting a design running. It's

NOTE 1: For recommended magnetics supplier, see page 13.

 LTC and LT are registered trademarks of Linear Technology Corporation.

Application Note 25

much easier to work with, and get enthusiastic about, a functional circuit than the smoking carcass of a devastated breadboard. If standard inductor characteristics aren't optimal, it's easier to see the evidence on a 'scope than to guess why you don't see anything.

Additionally, once the circuit is running, an optimized version of the standard product can be supplied by the inductor manufacturer. It's generally easier for the inductor manufacturer to modify its standard product than to start from scratch. The process of communicating and translating circuit performance requirements into inductor construction details is tricky. Using standard product as a starting point accelerates the dialogue, minimizing the number of iterations required for satisfactory results. Often, the standard product suffices for the purpose and no further effort is required.

Strictly speaking, it makes more sense to design the inductor to meet circuit requirements than to fashion a circuit around a standard inductor. Deliberately ignoring this consideration considerably complicated the author's work, but hopefully will simplify the reader's (such is the lot of an application note writer's life). Those interested in inductor design theory are commended to LTC Application Note AN-19, "LT[®]1070 Design Manual."

A final aid in achieving my new outlook on switchers was the LT1070 family. In terms of circuit construction and ease of use they really are superior switching regulator ICs. A 75V, 5A (LT1070HV) on-chip power switch, complete control loop, oscillator and only 5 pins eliminate a lot of the ambiguity of other devices. Internal details and operating features of the LT1070 family are detailed in Appendix A, "Physiology of the LT1070."

Basic Flyback Regulator

Figure 1 shows a basic flyback regulator using the LT1070. It converts a 5V input to a 12V output. Figure 2 shows the voltage (Trace A) and the current (Trace B) waveforms at the V_{SWITCH} pin. The V_{SW} output is the collector of a common emitter NPN, so current flows when it is low. Current is pulled through the 100 μ H inductor and controlled to a value of which forces the 12V output to be constant. The circuit's 40kHz repetition rate is set by the LT1070's internal oscillator. During the time V_{SW} is low, current flow through the inductor causes a magnetic field

to be induced into the area around the inductor. The amount of energy stored in this field is a function of the current level, how long current flows, the characteristics of the inductor and its core material. It is often useful to think of the inductor as a bucket and analogize current flow as water pouring into it. The ultimate limit on energy storage is set by the bucket's capacity, corresponding to the inductor's saturation limitations. The amount of energy that can be put into an inductor in a given time is limited by the applied voltage and the inductance. The amount of energy that can be stored without saturating the inductor is limited by the core characteristics. Size, core material, operating frequency, voltage and current influence inductor design.




Figure 1. Flyback-Type Regulator




Figure 2. Flyback Regulator's Waveforms at 7W Loading

If the inductor is enclosed in a feedback-enforced loop, such as Figure 1, the energy put into it will be controlled to meet circuit output demands. Figure 3 shows what happens when output demand doubles. In this case duty




Figure 3. Flyback Regulator's Waveforms at 14W Loading

cycle doesn't change much but current doubles. This requires the inductor to store more energy. If it couldn't meet the storage requirement, e.g., it saturated and could not hold any more magnetic flux, it would cease to look inductive. If this point is reached, current flow is limited only by the resistance of the wire and rapidly builds to excessive and destructive values. This behavior is exactly the opposite of a capacitor, where current diminishes upon entering saturation. Capacitors can maintain energy storage with no current flowing; inductors cannot. See Appendix C, "A Checklist for Switching Regulator Designs," for details.

At the end of each inductor charge cycle, current flow in the inductor decays, and the magnetic field around it abruptly collapses. The V_{SW} pin is seen to rise rapidly to a voltage higher than the 5V input. This flyback action gives the regulator its voltage boost characteristics and its name. The boost characteristic is caused by the collapsing magnetic field's lines of flux cutting across the inductor's conductive wire turns. This satisfies the basic requirement for generation of a current in (and hence, a voltage across) a conductor. This moving magnetic field deposits energy into the wire in proportion to how much was stored in the core during the current charge cycle. It is worth noting that the operating characteristics shown here are similar to the Kettering ignition system used in automobiles, explaining why spark occurs when the points open.²

In this circuit the flyback is seen to clamp to a level just above the output voltage. This is so because the flyback pulse is steered through the Schottky diode to the output. The 470µF capacitor integrates the repetitive flyback events to DC, providing the circuit's output. The feedback pin (FB)

samples this output via the 10.7k to 1.24k divider. The LT1070 compares the feedback pin voltage to its internal 1.24V reference and controls the V_{SW} pin's duty cycle and current, closing a loop. Since the LT1070 is trying to force its feedback pin to 1.24V, output voltage may be set by varying the 10.7k or 1.24k values.

All feedback loops require some form of stability compensation (see the appended section of LTC Application Note AN-18, "The Oscillation Problem—Frequency Compensation Without Tears," for general discussion). The LT1070 is no exception. Its voltage gain characteristics, combined with the substantial phase shift of the circuit's sampled energy delivery, ensure oscillation if uncompensated. While the large output capacitor smooths the output to DC, it also teams up with the sampled energy coming into it to create phase shift. To complicate matters, the load, which may vary, also influences phase characteristics. The regulator can only source into the output capacitor. The load determines the sink time constant, influencing phase performance and overall stability.

The LT1070's internals have been designed with all this in mind and compensation is usually fairly simple. In this case the 1k to 1µF combination at the compensation pin (V_C) rolls off the circuit, providing stable compensation for all operating conditions (see Appendix B, "Frequency Compensation," for details and suggestions on achieving stability in switching regulator loops).

As innocent as Figure 1 appears, it's not too difficult to get into odd and seemingly inexplicable problems. Note that the ground connection appears at the ground pin, as opposed to its customary location at the bottom of the diagram. This is deliberate and the supply and load return connections should be made there. The high speed, high current returns from the output transistor's emitter (the "other end" of the V_{SW} pin) should not be allowed to mix with the small currents of the output divider or the V_C pin.

Such mixing can promote poor regulation, unstable operation or outright oscillation. Similarly, the 22µF bypass capacitor ensures clean local power at the LT1070, even during the fast, high current drain periods when V_{SW} comes on. It should have good high frequency characteristics (tantalum or aluminum paralleled by a disc ceramic type). More discussion of these considerations appears in Appendix C.

NOTE 2: Back when giants walked the earth, Real Cars used ignition points.

Application Note 25




Figure 4. Nonisolated -48V to 5V Regulator

-48V to 5V Telecom Flyback Regulator

Figure 4's circuit is operationally similar to Figure 1 but is intended for telecom applications. The raw telecom supply is nominally -48V but can vary from -40V to -60V . This range of voltages is acceptable to the V_{SW} pin but protection is required for the V_{IN} pin ($V_{\text{MAX}} = 60\text{V}$). Q1 and the 30V zener diode serve this purpose, dropping V_{IN} 's voltage to acceptable levels under all line conditions.

Here, the “top” of the inductor is at ground and the LT1070’s ground pin at -48V . The feedback pin senses with respect to the ground pin, so a level shift is required for the 5V output. Q2 serves this purpose, introducing only $-2\text{mV}/^\circ\text{C}$ drift. This is normally not objectionable in a logic supply, but can be compensated with the optional appropriately scaled diode-resistor shown.

Frequency compensation is similar to Figure 1, although a low ESR (equivalent series resistance) capacitor gives less phase shift, permitting faster loop response with the reduced compensation time constant. The 68V zener is a type designed to clamp and absorb excessive line

transients which might otherwise damage the LT1070 (V_{SW} maximum voltage is 75V).

Figure 5 shows operating waveforms at the V_{SW} pin. Trace A is the voltage and Trace B the current. Switching characteristics are fast and clean. The ripples in the current trace are due to nonoptimal breadboard layout (ground as I say, not as I do). Inductor ringing on turn-off (Trace A) is characteristic of flyback configurations.




Figure 5. Nonisolated Regulator's Waveforms




Figure 6. Fully Isolated -48V to 5V Regulator

Fully-Isolated Telecom Flyback Regulator

Figure 6's circuit is another telecom regulator. Although it looks more complex, it's really a closely related extension of the previous flyback circuits. The fundamental difference is that the output is fully galvanically isolated from the input, often a requirement in equipment. This necessitates a transformer instead of a simple 2-terminal inductor. It also requires output feedback information to be transmitted to the regulator across a nonconducting path. The transformer complicates the circuit's start-up and switching characteristics while the isolated feedback requires attention to frequency compensation.

In this circuit the V_{IN} pin receives power from a transformer winding. This winding cannot supply power at start-up because the circuit is nonfunctional. Q1 through Q4 address this issue. When power is applied, Q5 cannot conduct because the LT1071 is unpowered. Q1 zener-

connected Q2 and Q3 are off. Under these conditions Q4 is on, pulling the V_C pin down and strobing off the LT1071. The potential at Q1's emitter slowly rises as the 10k-100µF combination charges. When Q1's emitter rises high enough, it turns on. Zener-connected Q2 conducts when the voltage across it is about 7V, biasing Q3 on. Q1 sees regenerative feedback, turning Q3 on harder. Q3's turn-on cuts off Q4, allowing the V_C pin to rise and biasing up the LT1071. The rate of rise is limited by the 10µF diode combination at the V_C pin. This network forces the V_C pin to come up slowly, providing a soft-start characteristic (the 100Ω diode string discharges the 10µF capacitor when circuit input power is removed). Because of this sequence, the LT1071 cannot start up the circuit until the V_{IN} potential is well established. This prevents start-up at "starved" or unstable V_{IN} voltages which could cause erratic or destructive modes. When start-up does occur, the transformer feeds the V_{IN} pin with DC via the MUR120 diode.

Application Note 25

The 50Ω resistor combines with the $100\mu\text{F}$ capacitor to give good ripple and transient filtering. This voltage is ample to run the LT1071 and reduces the current through the 10k resistor, saving power. Q1, Q2 and Q3 remain on, biasing Q4 to allow LT1071 operation.

In the previous flyback circuits, the V_{SW} pin drove the inductor directly. Here, a power MOSFET is interposed between the V_{SW} pin and the inductor. In this arrangement the inductor is a transformer and its flyback characteristics are different from a simple 2-terminal inductor. For the simple inductor, the flyback energy was clamped by and dumped directly into the output capacitor. Excessive voltages did not occur. In the transformer case, all the flyback energy does not end up in the output capacitor. Substantial flyback voltage spikes ($>100\text{V}$) appear across the transformer primary when the LT1071 driven MOSFET turns off.

Several measures prevent these spikes from destroying the circuit. The $0.47\mu\text{F}$ - 2k -diode combination, a damper network, conducts during the flyback event. This loads the transformer primary, minimizing flyback amplitude. The damper values are selected empirically, with the trade-off being power dissipation in them. Very low values markedly reduce flyback potentials but cause excessive dissipation. High values permit low dissipation but allow excessive flyback voltages. The damper values should be selected under fully-loaded output conditions because flyback energy is proportionate to transformer power levels. Appendix C contains additional information on damper network considerations.

Even with the damper network, the flyback voltage is too high for the LT1071 output transistor. Q5 prevents the LT1071 from seeing the high voltage. It is connected in series with the LT1071's output transistor. This connection, sometimes called a cascode, lets Q5 stand off the high voltage and the LT1071 operates well within its breakdown limits. Development and testing of this configuration is detailed in Appendix D. Q5 has large parasitic capacitances associated with all terminals. During

switching, these capacitances can cause excessive transient voltages to appear. The 18V zener diode insures against gate-source breakdown ($V_{\text{GS MAX}} = 20\text{V}$) and the diode clamps the V_{SW} pin to the V_{IN} potential. Mention of these considerations appears in Appendix C.

The transformer's rectified and filtered secondary produces the 5V output. This output is galvanically isolated from the circuit's input. To preserve this desired feature, the feedback path must also be galvanically isolated. A1, the optoisolator and their associated components serve this function. A1, powered by the 5V output, compares a resistively-sampled portion of the output with the LT1004 1.2V reference. Operating at a gain of 200, it drives the optoisolator's LED. The optoisolator's output transistor biases the LT1071's V_{C} pin, closing a regulation loop. The feedback amplifier inside the LT1071 is essentially bypassed by the A1 optoisolator combination and is not used. Normally, the optoisolator's drift transmission characteristics over time and temperature would result in unstable feedback. Here, A1's gain is placed ahead of the optoisolator. This attenuates these uncertainties, providing a stable loop. This approach is not too different from inside-the-loop booster transistors and buffers used with op amps. Both schemes rely on the op amp's gain to eliminate uncertainties and drifts. Returning the optoisolator to V_{REF} instead of ground forces the op amp to bias well above ground, minimizing saturation effects during output transients.

Frequency compensation is somewhat more involved in this circuit than the previous examples. A1 is rolled off by the $0.1\mu\text{F}$ unit. This keeps gain low at high frequency, preventing amplified ripple and noise from being fed back to the LT1071. Local compensation at the LT1071 V_{C} pin stabilizes the loop. The 100Ω resistor at the 5V output, a deliberate sink path, allows loop stability at light or no load. Appendix B discusses frequency compensation.

Additional transformer secondary windings could be added if desired. The input zener clips transient voltages.

Circuit waveforms appear in Figure 7. Trace A is Q5's drain voltage and Trace B the drain current. Trace A shows that the MOSFET sees about 100V due to flyback effects, but this is well within its rating. The ringing on turn-off is normal and is similar to the waveform observed in Figure 4's circuit. Trace B shows that the current flow is fast, clean and controlled. Figure 8 shows transient response for a 1A step on a 2.5A output. When Trace A goes high the step occurs. Trace B shows that output sag is corrected in about 8ms. When Trace A returns low the 1A load is removed and recovery is similar to the positive step. Broadband output noise, about 75mV_{p-p}, may be reduced with the optional filter shown.




Figure 7. Fully Isolated Regulator's Waveforms




Figure 8. Fully Isolated Regulator's Transient Response for a 1A Change on a 2.5A Load

100W Off-Line Switching Regulator

One of the most desirable switching regulator circuits is also one of the most difficult to design. Figure 9's circuit has many similarities to the previous design but is powered directly from the 115V AC line. This off-line operation is desirable because it eliminates large, heavy and inefficient 60Hz magnetics and filter capacitors. The circuit provides an isolated 5V, 20A output as well as isolated $\pm 12V$, 1A outputs. Additional features include operation over a 90V AC to 140V AC input range, AC line surge suppression, soft-starting and loop stability under all conditions. Efficiency exceeds 75%.

BEFORE PROCEEDING ANY FURTHER, THE READER IS WARNED THAT CAUTION MUST BE USED IN THE CONSTRUCTION, TESTING AND USE OF THIS CIRCUIT. HIGH VOLTAGE, AC LINE-CONNECTED POTENTIALS ARE PRESENT IN THIS CIRCUIT. EXTREME CAUTION MUST BE USED IN WORKING WITH AND MAKING CONNECTIONS TO THIS CIRCUIT. REPEAT: THIS CIRCUIT CONTAINS DANGEROUS, AC LINE-CONNECTED HIGH VOLTAGE POTENTIALS. USE CAUTION.

AC line power is rectified and filtered by the diode bridge—470 μ F combination. The MOV device provides surge suppression and the thermistor limits turn-on in-rush current. Start-up and soft-start circuitry is similar to Figure 6's circuit, with some changes necessitated by the higher input voltage. Erratic operation at extremely low AC line voltages (70V AC) is prevented by the 220k-1.24k divider. At very low AC line inputs, this divider forces the LT1071 feedback pin to a low state, shutting down the circuit. The high input voltage, typically 160V DC, means that the LT1071's internal current limit is set too high to protect the regulator if the circuit's output is shorted. Q6 and its associated components provide about 2A limiting. The LT1071's GND pin current flows through the 0.3 Ω resistor, turning on Q6 if current is too high. The 22k-50pF RC filters noise, preventing erratic Q6 operation.

Application Note 25

Q5, a power MOSFET, is cascoded with the LT1071 for high voltage switching. Circuit topology is similar to Figure 6, with Q5's voltage breakdown increased to 500V.

Additionally, the 50Ω resistor combines with the gate capacitance to slightly slow Q5's transitions, reducing high frequency harmonics. This measure eases layout considerations. The transformer's damper network borrows from Figure 6, with values reestablished for this circuit.

The A1-optocoupler-enforced feedback loop preserves the transformer's galvanic isolation, allowing the regulator output to be ground referenced. The feedback loop is also similar to Figure 6. Compensation values at A1 and the

LT1071 have changed, reflecting this circuit's different gain-phase characteristics.

ALL WAVEFORM PHOTOGRAPHS WERE TAKEN WITH AN ISOLATION TRANSFORMER CONNECTED BETWEEN THE CIRCUIT'S 90V AC-140V AC INPUT AND THE AC LINE. USERS AND CONSTRUCTORS OF THIS CIRCUIT MUST OBSERVE THIS PRECAUTION WHEN CONNECTING TEST EQUIPMENT TO THE CIRCUIT TO AVOID ELECTRIC SHOCK. REPEAT: AN ISOLATION TRANSFORMER MUST BE CONNECTED BETWEEN FIGURE 9'S CIRCUIT AND THE AC LINE IF ANY TEST EQUIPMENT IS TO BE CONNECTED.




Figure 9. 100W Off-Line Switching Regulator
DANGER! Lethal Potentials Present—See Text

Figure 10 shows circuit waveforms at 15A output. Trace A, Q5's drain, shows the flyback pulse being damped below 300V (for a discussion of the procedures used to design the damper network and other design techniques in this circuit, see Appendix D, "Evolution of a Switching Regulator Design"). Trace B, the LT1071's V_{SW} pin, stays well within its voltage rating, despite Q5's high voltage switching. Trace C, Q5's drain current, shows that transformer current is well controlled with no saturation effects. Trace D, damper network current, is active when Q5 goes off.

Figure 11 is a time and amplitude expansion of Q5's drain (Trace A) and transformer primary current (Trace B). Switching is clean, with residual noise due to nonideal transformer behavior. The damper network clamps the

flyback pulse well below Q5's 500V rating and the transformer rings off after the flyback interval. The noise on the current pulse, due to resonances in the transformer, has no significant effect on circuit operation.

Figure 12 shows output noise with the optional LC filter in use. Without the filter, noise is about 150mV. Superimposed, residual 120Hz modulation accounts for trace thickening at the peaks and could be eliminated by increasing the 470 μ F value.

Figure 13 shows transient response performance. When Trace A goes high, a 5A transient is added to a 10A steady-state load. Recovery amplitude is low and clean with a first order response. When Trace A goes low, the transient load is removed with similar results.




Figure 10. Off-Line Switcher's Waveforms
DANGER! Take This Measurement *Only* With an Isolation Transformer in Use—See Text




Figure 11. Detail of Off-Line Switcher's Transformer Primary Voltage and Current Waveforms
DANGER! Take This Measurement *Only* With an Isolation Transformer in Use—See Text




Figure 12. Figure 9's Output Ripple at 10A Output with the Optional LC Filter Added—Without the Filter, Ripple Increases to About 150mV_{p-p}




Figure 13. Figure 9's Circuit Responding to a 5A Change on a 10A Output

Application Note 25

Figure 14 shows response for shifts in the line. When Trace A is high, the AC line is at 140V AC. Line voltage drops to 90V AC with Trace A low. Trace B, the regulator's AC-coupled output, shows a clean recovery with small amplitude error. The ripples in the waveform, 120Hz input residue, could be reduced by increasing the 470 μ F capacitor.

Figure 15 shows the 5V output at start-up into a 20A load. Response is slightly underdamped and can be modified by adjusting the frequency compensation. The compensation shown in Figure 9 is a good compromise between transient response and turn-on characteristics. The delay on turn-on and the controlled rise time are due to the slow-start circuitry.

Figure 16 plots regulator efficiency. As would be expected, efficiency is best at high currents, where static losses are a small percentage of output power.

Switch-Controlled Motor Speed Controller

Voltage regulators are not the only switching power circuits. Figure 17 shows a motor speed regulator. The LT1070 provides simplicity and switch-mode control efficiency. Although this circuit controls a motor, it shares many considerations common to voltage regulators. When power is applied, the tachometer output is zero and the feedback pin (FB) is also at zero. This causes the LT1070 to begin pulsing its V_{SW} pin at maximum duty cycle. The motor turns, forcing tachometer output. When the FB pin




Figure 14. Figure 9 Responds to a 90V AC—140V AC Line Change—Loading is 10A—120Hz Residue in Output Could be Reduced by Increasing the 470 μ F Input Filter





Figure 15. Start-Up for Figure 9 at 20A Loading—The 10 μ F Capacitor at the LT1070's V_C Pin Produces the Slow-Start Characteristic. If the Small Overshoot is Objectionable, Modified Frequency Compensation Can Eliminate it at Some Cost to Transient Response


arrives at the LT1070's internal voltage reference value (1.24V), the loop stabilizes. Speed is adjustable with the 25k potentiometer in the feedback string. The MUR120 damps the motor's flyback spike. The characteristics of the motor specified permit no current limiting in series with the diode. Other motors might require this and

damper network optimization should be done for any specific unit. Similarly, frequency compensation values will vary with different motor types. The diode at the tachometer output prevents transient reverse voltages due to tachometer commutator switching.



AN-25 F16

Figure 16. Figure 9's Efficiency vs Operating Point



AN-25 F17

Figure 17. A Simple Motor-Tachometer Servo Loop

Application Note 25

Switch-Controlled Peltier 0°C Reference

Figure 18 is another switch-mode control circuit. Here, the LT1070 controls power to a Peltier cooler, providing a 0°C temperature reference for transducer calibration.

A platinum RTD is thermally mated to the Peltier cooler. The RTD combines with a bridge network to give a differential output. A1 provides maximum bridge drive without introducing significant heating in the RTD. The LTC[®]1043 switched capacitor network converts this output to a single-ended signal at A2. A2, operating at a gain of 400,

biases the LT1070's V_C pin. This closes a control loop around the Peltier cooler, forcing its temperature low enough to balance the bridge. The 0°C trim adjusts the servo point to precisely 0°C. A standard RTD should monitor Peltier temperature when making this trim. Alternatively, the sensor specified should be supplied with a certified 0°C resistance. With the RTD and Peltier cooler tightly mated, stability is excellent. Figure 19, a plot of stability over hours in a 25°C $\pm$ 3°C ambient, shows a 0.15°C baseline.




Figure 18. A Peltier-Cooled Switched-Mode 0°C Reference




Figure 19. Stability of Figure 18's Circuit Over Many Hours with a $25^{\circ}\text{C} \pm 3^{\circ}\text{C}$ Ambient

Acknowledgments

The author acknowledges Carl Nelson's abundance of commentary, some of which was useful, during preparation of this work. Bob Dobkin's thoughts and patience are also appreciated. Ron Young made significant contributions towards Figure 6's circuit. Bill McColey and other members of the Engineering staff of Pulse Engineering, Inc.³, supplied invaluable insight and assistance on magnetics issues. As usual, our customers' requests and requirements provided the most valuable source of guidance, and they are due a special thanks.

NOTE 3: P.O. Box 12235, San Diego, CA 92112 (619/268-2400)

APPENDIX A

Physiology of the LT1070

The LT1070 is a current-mode switcher. This means that switch duty cycle is directly controlled by switch current rather than by output voltage. Referring to Figure A1, the switch is turned on at the start of each oscillator cycle. It is turned off when switch current reaches a predetermined level. Control of output voltage is obtained by using the output of a voltage-sensing error amplifier to set current trip level. This technique has several advantages. First, it has immediate response to input voltage variations, unlike ordinary switchers which have notoriously poor line transient response. Second, it reduces the 90° phase shift at mid-frequencies in the energy storage inductor. This greatly simplifies closed loop frequency compensation under widely varying input voltage or output load conditions. Finally, it allows simple pulse-by-pulse current limiting to provide maximum switch protection under output overload or short conditions. A low dropout internal regulator provides a 2.3V supply for all internal circuitry on the LT1070. This low dropout design allows input voltage to vary from 3V to 6V with virtually no change in device performance. A 40kHz oscillator is the basic clock for all internal timing. It turns on the output switch via the logic and driver circuitry. Special adaptive antisat circuitry detects onset of saturation in the power switch and adjusts driver current instantaneously to limit switch saturation. This minimizes driver dissipation and provides very rapid turn-off of the switch.




Figure A1. LT1070 Internal Details

Application Note 25

A 1.2V bandgap reference biases the positive input of the error amplifier. The negative input is brought out for output voltage sensing. This feedback pin has a second function; when pulled low with an external resistor, it programs the LT1070 to disconnect the main error amplifier output and connects the output of the flyback amplifier to the comparator input. The LT1070 will then regulate the value of the flyback pulse with respect to the supply voltage. This flyback pulse is directly proportional to output voltage in the traditional transformer-coupled flyback topology regulator. By regulating the amplitude of the flyback pulse the output voltage can be regulated with no direct connection between input and output. The output is fully floating up to the breakdown voltage of the transformer windings. Multiple floating outputs are easily obtained with additional windings. A special delay network inside the LT1070 ignores the leakage inductance spike at the leading edge of the flyback pulse to improve output regulation.

The error signal developed at the comparator input is brought out externally. This pin (V_C) has four different functions. It is used for frequency compensation, current limit adjustment, soft-starting and total regulator shutdown. During normal regulator operation this pin sits at a voltage between 0.9V (low output current) and 2.0V (high output current). The error amplifiers are current output (gm) types, so this voltage can be externally clamped for adjusting current limit. Likewise, a capacitor-coupled external clamp will provide soft-start. Switch duty cycle goes to zero if the V_C pin is pulled to ground through a diode, placing the LT1070 in an idle mode. Pulling the V_C pin below 0.15V causes total regulator shutdown with only 50 μ A supply current for shutdown circuitry biasing. For more details, see Linear Technology Application Note AN-19, pages 4-8.

APPENDIX B

Frequency Compensation

Although the architecture of the LT1070 is simple enough to lend itself to a mathematical approach to frequency compensation, the added complications of input and/or output filters, unknown capacitor ESR, and gross operating point changes with input voltage and load current variations all suggest a more practical empirical method. Many hours spent on breadboards have shown that the simplest way to optimize the frequency compensation of the LT1070 is to use transient response techniques and an R/C box to quickly iterate toward the final compensation network.

There are many ways to inject a transient signal into a switching regulator, but the suggested method is to use an AC-coupled output load variation. This technique avoids problems of injection point loading and is general to all switching topologies. The only variation required may be an amplitude adjustment to maintain small signal conditions with adequate signal strength. Figure B1 shows the setup.

A function of generator with 50 Ω output impedance is coupled through a 50 Ω /1000 μ F series RC network to the regulator output. Generator frequency is noncritical. A good starting point is 50Hz. Lower frequencies may cause a blinking scope display which is annoying to work with. Higher frequencies may not allow sufficient settling time for the output transient. Amplitude of the generator output is typically set at 5V_{P-P} to generate a 100mA_{P-P} load variation.

For lightly loaded output ($I_{OUT} < 100$ mA), this level may be too high for small signal response. If the positive and negative transition settling waveforms are significantly different, amplitude should be reduced. Actual amplitude is not particularly important because it is the shape of the resulting regulator output waveform that indicates loop stability.

A 2-pole oscilloscope filter with $f = 10$ kHz is used to block switching frequencies. Regulators without added LC output filters have switching frequency signals at their

outputs which may have much higher amplitude than the low frequency settling waveform to be studied. The filter frequency is high enough to pass the settling waveform with no distortion.

Oscilloscope and generator connections should be made exactly as shown to prevent ground loop errors. The oscilloscope is synced by connecting the channel B probe to the generator output, with the ground clip of the second probe connected to exactly the same place as the channel A ground. The standard 50Ω BNC sync output of the generator should not be used because of ground loop errors. It may also be necessary to isolate either the generator or oscilloscope from its third wire (earth ground) connection in the power plug to prevent ground loop errors in the 'scope display. These ground loop errors are checked by connecting the channel A probe tip to exactly

the same point as the probe ground clip. Any reading on channel A indicates a ground loop problem.

Once the proper setup is made, finding the optimum values for the frequency compensation network is fairly straightforward. Initially, C2 is made large ($\geq 2\mu\text{F}$), and R3 is made small ($\approx 1\text{k}\Omega$). This nearly always ensures that the regulator will be stable enough to start iteration. Now, if the regulator output waveform is single-pole overdamped (see the waveforms in Figure B2), the value of C2 is reduced in steps of about 2:1 until the response becomes slightly underdamped. Next, R3 is increased in steps of 2;1 to introduce a loop "zero." This will normally improve damping and allow the value of C2 to be further reduced. Shifting back and forth between R3 and C2 variations will now allow one to quickly find optimum values.




Figure B1. Testing Loop Stability




Figure B2. Output Transient Response

Application Note 25

If the regulator response is underdamped with the initial large value of C, R should be increased immediately before larger values of C are tried. This will normally bring about the overdamped starting condition for further iteration.

Just what is meant by “optimum values” for R3 and C2? This normally means the smallest value for C2 and the largest value for R3, which still guarantee no loop oscillations, and which result in loop settling that is as rapid as possible. The reason for this approach is that it minimizes the variations in output voltage caused by input ripple voltage and output load transients. A switching regulator which is grossly overdamped will never oscillate but it may have unacceptably large output transients following sudden changes in input voltage or output loading. It may also suffer from excessive overshoot problems on start-up or short-circuit recovery.

To guarantee acceptable loop stability under all conditions, the initial values chosen for R3 and C2 should be checked under all conditions of input voltage and load current. The simplest way to accomplish this is to apply load currents of minimum, maximum, and several points in between. At each load current, input voltage is varied

from minimum to maximum while observing the settling waveform. The additional time spent “worst-casing” in this manner is definitely necessary. Switching regulators, unlike linear regulators, have large shifts in loop gain and phase with operating conditions. If large temperature variations are expected for the regulator, stability checks should also be done at the temperature extremes. There can be significant temperature variations in several key component parameters which affect stability—in particular, input and output capacitor values and their ESRs and inductor permeability. The LT1070 parametric variations also need some consideration. Those which affect loop stability are error amplifier gm, and the transfer function of V_C pin voltage versus switch current (listed as a transconductance under electrical specifications.) For modest temperature variations, conservative overdamping under worst-case temperature conditions is usually sufficient to guarantee adequate stability at all temperatures.

If external amplifiers or other active devices are included in the loop (e.g., Figures 6 and 9), their effects must be included in stabilizing the loop. LTC Application Note AN-18, pages 12-15, provides commentary that may be useful in these situations.

APPENDIX C

A Checklist for Switching Regulator Designs

1. The most common problem area in switching designs is the inductor and the most common difficulty is saturation. An inductor is saturated when it cannot hold any more magnetic flux. As an inductor arrives at saturation it begins to look more resistive and less inductive. Under these conditions the current flow through it is limited only by its DC copper resistance and the source capacity. This is why saturation often results in destructive failures. Figure C1 demonstrates saturation effects. The pulse generator drives Q1, forcing current into the inductor. The diode and RC combination form a typical load. Figure C2 shows results. The voltage at Q1's collector falls when it turns on (Trace A is pulse generator output, Trace B is Q1's collector). Trace C, the inductor current, ramps in controlled fashion. When Q1 goes off, current falls and the induc-

tor rings off. In Figure C3, drive pulse width is longer, allowing more inductor current buildup. This requires the inductor to store more magnetic flux. Its ramp waveform is clean and controlled, indicating that it has the necessary capacity. Figure C4 brings some unpleasant surprises. Drive pulse width has been increased. Now, the inductor current departs from its ramp characteristic into a nonlinear slope. The nonlinear behavior starts between the third and fourth vertical divisions. This curve shows a rapidly increasing current characteristic. These conditions indicate that the inductor is entering saturation. If pulse width is increased much more, the current will rise to destructive levels. It is worth noting that some inductors saturate much more abruptly than this case.




Figure C1. Inductor Saturation Test Circuit




Figure C2. Normal Inductor Operation




Figure C3. Normal Inductor Operation at Increased Current




Figure C4. Inductor Being Driven into Saturation

Application Note 25

2. Always consider inductive flyback effects. Are semiconductor breakdown ratings adequate to withstand them? Is a snubber (damper) network required? Consider all possible voltages and current paths, including the transient ones via semiconductor junction capacitances, to avoid evil problems.
3. Think about requirements in capacitors. All operating conditions should be accounted for. Voltage rating is the most obvious consideration, but remember to plan for the effects of equivalent series resistance (ESR) and inductance. These specifications can have significant impact on circuit performance. In particular, an output capacitor with high ESR can make loop compensation difficult.
4. Layout is vital. Don't mix signal, frequency compensation, and feedback returns with high current returns. Arrange the grounding scheme for the best compromise between AC and DC performance. In many cases, a ground plane may help. Account for possible effects of stray inductor-generated flux on other components and plan layout accordingly.
5. Semiconductor breakdown ratings must be thought through. Account for all conditions. Transient events usually cause the most trouble, introducing stresses that are often hard to predict. Things to watch for include effects of feedthrough via semiconductor junction capacitances (note the clamping of Q5's gate in Figures 6 and 9). Such capacitances can allow excessive voltages to occur for brief durations at what is normally a low voltage node. Study the data sheet breakdown, current capacity, and switching speed ratings carefully. Were these specifications written under the same conditions that your circuit is using the device in? If in doubt, consult the manufacturer.

"Simple" diodes furnish a good example of how carefully semiconductor operating conditions must be considered in switching regulators. Switching diodes have two important transient characteristics—reverse recovery time and forward turn-on time. Reverse recovery time occurs because the diode stores charge during its forward

conducting cycle. This stored charge causes the diode to act as a low impedance conductive element for a short period of time after reverse drive is applied. Reverse recovery time is measured by forward biasing the diode with a specified current, then forcing a second specified current backwards through the diode. The time required for the diode to change from a reverse conducting state to its normal reverse nonconducting state is reverse recovery time. Hard turn-off diodes switch abruptly from one state to the other following reverse recovery time. They therefore dissipate very little power even with moderate reverse recovery times. Soft turn-off diodes have a gradual turn-off characteristic that can cause considerable diode dissipation during the turn-off interval. Figure C5 shows typical current and voltage waveforms for several commercial diode types used in an LT1070 flyback converter with $V_{IN} = 10V$, $V_{OUT} = 20V$, $2A$.




Figure C5. Diode Turn-Off Characteristics

Long reverse recovery times can cause significant extra heating in the diode or the LT1070 switch. Total power dissipated is given by:

$$P_{trr} = V \cdot f \cdot t_{RR} \cdot I_F$$

V = reverse diode voltage

f = LT1070 switching frequency

t_{RR} = reverse recovery time

I_F = diode forward current just prior to turn-off

With the circuit mentioned, I_F is 4A, V = 20V, and f = 40kHz. Note that diode on current is twice output current for this particular boost configuration. A diode with t_{rr} = 300ns creates a power loss of:

$$P_{trr} = (20)(40 \cdot 10^3)(300 \cdot 10^{-6})(4) = 0.96W$$

If this same diode had a forward voltage of 0.8V at 4A, its forward loss would be 2A (average current) times 0.8V equals 1.6W. Reverse recovery losses in this example are nearly as large as forward losses. It is important to realize however, that reverse losses may not necessarily increase diode dissipation significantly. A hard turn-off diode will shift much of the power dissipation to the LT1070 switch, which will undergo a high current and high voltage condition during the duration of reverse recovery time. This has not been shown to be harmful to the LT1070, but the power loss remains.

Diode turn-on time can potentially be more harmful than reverse turn-off. It is normally assumed that the output diode clamps to the output voltage and prevents the inductor or transformer connection from rising higher than the output. A diode that turns on slowly can have a very high forward voltage for the duration of turn-on time. The problem is that this increased voltage appears across the LT1070 switch. A 20V turn-on spike superimposed on a 40V flyback mode output pushes switch voltage perilously close to the 65V limit. The graphs in Figure C6

show diode turn-on spikes for three common diode types—fast, ultrafast, and Schottky. The height of the spike will be dependent on rate of rise of current and the final current value, but these graphs emphasize the need for fast turn-on characteristics in applications which push the limits of switch voltage.

Fast diodes can be useless if the stray inductance is high in the diode, output capacitor or LT1070 loop. 20-gauge hook-up wire has 30nH/inch inductance. The current fall time of the LT1070 switch is 10^8 A/sec. This generates a voltage of $(10^8)(30 \cdot 10^{-9}) = 3V$ per inch in stray wiring. Keep the diode, capacitor and LT1070 ground/switch lead lengths *short!*




Figure C6. Diode Turn-On Spike

Figure D3 shows the test circuit rearranged to accommodate the transformer load. The transformer replaces the resistor. Its terminated secondary allows it to present a significant load. The fixed 160V supply has been replaced with a 0V to 200V unit, permitting voltage to be slowly and cautiously increased¹. The 350V transistor is replaced with a 1000V unit, in preparation for inductive events. Figure D4 shows waveforms. As expected, the inductive flyback (Trace A) is significant, even at low supply voltage ($V_{SUPPLY} = 60V$ in this photo).

Trace C, the drain current, rises with a characteristic indicating the inductive load. Trace B, the source voltage, is of greater concern. The flyback event, feeding through the MOSFET's capacitances, causes the source (and gate) to rise above nominal clamped value. At the higher supply voltages planned, this could cause excessive gate-source voltages with resultant device destruction. Because of this, the zener diode in dashed lines is installed, clamping gate-source voltages to safe values. This component appears in Figure 9's final design. With this correction, behavior at higher supply voltages may be investigated.




Figure D3. Test Circuit for MOSFET-LT1071 Cascode with Transformer Load




Figure D4. MOSFET-LT1071 Cascode Switching the Transformer Primary—Secondary Load is 0.2Ω

NOTE 1: "For fools rush in where angels fear to tread"—An Essay on Criticism, A. Pope.

Application Note 25

Figure D5 shows the MOSFET drain at $V_{\text{SUPPLY}} = 160\text{V}$. The load draws about 2.5A. Flyback voltage rises to 400V. At 5A loading this voltage approaches 500V (Figure D6), while a 10A load (Figure D7) forces almost 900V flyback. In actual regulator operation, supply voltages, switch on-time and output current can go higher, meaning flyback

potentials will exceed 1000V. This graphically mandates the need for a damper network. A simple reverse-biased diode or zener clipper will work, but will suffer from excessive dissipation. The network shown in Figure 9 is a good compromise between dissipation and reasonable flyback voltages.




Figure D5. Undamped Regulator Flyback Pulse at 2.5A Output




Figure D6. Undamped Regulator Flyback Pulse at 5A Output




Figure D7. Undamped Regulator Flyback Pulse at 10A Output

Once the drive-flyback issues are settled, a feedback loop is closed around the transformer. This allows checking to see that loop stabilization is possible. Figure D8 diagrams the loop. In this configuration the regulator will function, but is unusable. The output is not galvanically isolated from the input, which ultimately must be directly AC line driven. After this loop has been successfully closed, the isolated version is tried (Figure D9). This introduces more

phase shift, but is also found to be stable with appropriate frequency compensation. Finally, the connection between the input and output common potentials is broken, achieving the desired galvanic isolation. The start-up, soft-start and current limit features are then added and optimized. Testing involves checking performance under various line and load conditions. Details on circuit operation are covered in the text associated with Figure 9.




Figure D8. Developmental Version of Off-Line Switching Regulator—No Isolation is Included and the Scheme is Solely Intended to Verify that a Loop Can be Closed Around the Transformer





Figure D9. Developmental Version of Off-Line Regulator with Isolation— the Circuit Verifies That Loop Stability is Achievable with the Added Phase Shift of A1 and the Opto-Isolator—Start-Up, Current Limit and Soft-Start Features Must be Added to Complete the Design

Application Note 25



(They told me I couldn't
leave the last page blank)

William

'87