2

Practical Selection of Topology

INTRODUCTION: THERE ARE HUNDREDS OF TOPOLOGIES!

Before you can begin any sort of design work on a converter, you have to select a topology. This is a really important task, as all other design selections depend on it: component selection, magnetics design, loop compensation, and so on; if the topology changes, these must change as well. So before getting started, it's always a good idea to spend some time carefully looking at the power supply's requirements and specifications to ensure that a proper topology is selected.

But how to choose? Some books on power supplies are nothing but compendiums of dozens of topologies, each with a few paragraphs describing the general idea of how the topology works, but little or nothing about the pros and cons of each, and certainly without guidance as to how to select one out of the many. Indeed, it has been shown in recent classification papers (see, e.g., Ref. 1) that resonant topologies alone number in the hundreds!

In this chapter, we're going to do it more practically. We're going to mention only the half-dozen or so topologies that are most commonly used in the low to medium power range, and clearly spell out their pros and cons. This book can't give absolute guidelines about which topology to use, because in fact you can make almost any of them work for a given application; but it will give strong opinions about which topologies not to use when, and the reasons why. In the first section on general considerations, we list the various criteria you need to consider when selecting a topology. The remainder of the chapter discusses the common topologies and some of their aspects vis-à-vis the criteria.

GENERAL CONSIDERATIONS

Step-Up or Step-Down

One of the very first things you need to think about to select a topology is whether the output voltage or voltages is (are) higher or lower than the input voltage, and whether this is true over the whole range of input voltages. For example, the buck converter can only step down the voltage; so the output voltage has to be less than the input at all times. (Details about the various types of converter mentioned in this section can be found in the sections of the chapter below.) If you have a 24V input that you want to step down to 15V, that's fine for a buck; but if the 24V actually has a range from 8V to 80V (as in MIL-STD-704A) then you can't use a buck, because you can't have 8V in and 15V out.

Practical Limits on Duty Cycle

Furthermore, there is a practical limit to how large or small a conversion ratio (output voltage divided by input voltage) can be achieved with a switching converter. First, the achievable duty cycle (definition: duty cycle = on time/switching period of a switch) for a converter has both a maximum and a minimum limit. In some topologies, you can't go above 50% duty cycle. In any case, commonly available PWM ICs often don't guarantee that they can reach duty cycles above about 85%. And in any case, many of them also don't function properly below about 5% duty cycle; which at reasonable switching frequencies is just as well, for you can't drive the gates of MOSFETs fast enough to get reasonable losses.

EXAMPLE

If your switching frequency is 250kHz, the period is 4μs. At a duty cycle of 10%, the on-time of the MOSFET is only 400ns, and if it takes 100ns to turn on the MOSFET and another 100ns to turn it off, almost all the period is eaten up in transitions, making for a lossy converter.

Practical Note Don't plan on running duty cycles outside the limits of approximately 10% minimum or 80% maximum (45% maximum for converters with a theoretical maximum of 50%), without taking special precautions (type of IC used, high current gate drives, etc.).

There is a way around the limitation in duty cycle just illustrated: by using a topology that has a transformer, you can achieve a greater conversion ratio by a factor of the turns ratio. However, there are limits even to this. If the turns ratio becomes extremely large, the gross mismatch in wire gauge between the primary and secondary makes the transformer difficult to wind.

Practical Note In general, transformers should have a maximum primary to secondary turns ratio of 10 : 1 or a minimum of 1 : 10. If you need to get really high voltages from a low voltage, or vice versa, you should think about either a two-stage converter or a voltage multiplier on the secondary.

How Many Outputs?

Closely connected with the question of duty cycle is the need to determine how many output voltages are to be generated. For example, if the answer is anything other than "one," a buck isn't suitable. Other practical limitations with some types of converter (concerning how many outputs should be planned on) are discussed below.

In the general sort of case, you may find that there are ways around such limits. For example, it may be possible to postregulate an output to generate another voltage. A common example might be a buck converter which produces a +5V output, and then uses a linear regulator (or even another switcher) with the +5V as input to generate +3.3V. The losses associated with this may be justifiable due, for example, to transient or noise requirements on the additional line.

In the worst case, it may make sense to design two separate converters, rather than trying to design and produce extremely complex magnetic pieces with large numbers of windings. Indeed, some of the worst converters the author has ever had to deal with (from the standpoint of producability and maintenance) have been multiwinding units, whose designer thought a few pennies could be saved by using a single PWM IC instead of two, and instead ended up spending dollars trying to make a very complicated transformer. The cost of magnetics should be considered up front, before any design is done, to avoid getting trapped with this problem.

Isolation

Another question that should be asked up front is whether primary to secondary isolation is required. There are all sorts of safety rules in the commercial world (as well as EMI questions, considered below) that may make isolation necessary. A typical example might be that the input has 500VAC applied to it relative to the output. But as soon as you know you need isolation, a number of topologies are immediately ruled out, that is, anything without a transformer (buck, nonisolated flyback, etc.).

EMI

Hopefully you've been lectured often enough, "Think about EMI from the start of the design, don't wait till the converter's already designed to start looking for Bandaids." Topology can have a lot to do with success in EMI. To start with the most basic aspect, if you have a nonisolated system, you have no common mode noise, since there is no third wire involved in the system! (EMI has a special chapter, Chapter 9, that explains these concepts in detail.) This makes filtering easier on you, the designer, if not easier overall.

Furthermore, some topologies are inherently more noisy than others. A distinction is to be made between topologies that disconnect the input from the converter during some portion of the period (and thus necessarily have discontinuous input current), and those that don't; the latter are easier to filter because the current has "less sharp edges." Among the latter, we distinguish those that operate in discontinuous mode from those that operate in continuous mode (this concept is discussed below), since the discontinuous operation also results in some portion of the period when the input current goes to zero; by the same reasoning, continuous mode will be easier to filter.

An example of a converter that disconnects the input is a buck, since when the switch is open, input current is zero. A nonisolated flyback always has the inductor connected to the input, but whether the input current is continuous depends on whether the flyback is being operated in continuous or discontinuous mode.

I recommend against using any of the topologies that claim they have no input ripple. Experience shows that they generally have very expensive magnetics.

Bipolar versus MOSFET versus ?

This question of what switch to use isn't directly related to topology selection, but should be considered up front also. The reason is that different types of devices have very different types of drive requirement; driving a bipolar transistor can be so hard that you will want to limit yourself to a single-switch topology. As of the date of writing, in the low to medium power range covered by this book, MOSFETs are used at least 90% of the time, both in commercial and military work. Indeed, except for special reasons, you should simply plan on using MOSFETs.

One of the special reasons is cost. For really high production quantities, a bipolar may still at times be cheaper than a MOSFET. However, a bipolar usually means a lower switching frequency than a MOSFET, and so the magnetics will be larger. Where does the cost advantage lie? You will have to do a detailed cost study to find this out.

You may also weigh the possibilities of a bipolar design for high input voltages, such as in 277V off-line conversion, or in a converter such as a push–pull, where you get double the input voltage, plus transients. You can get a 1500V bipolar, but the maximum MOSFET voltage is 1000V. Of course, for this you might consider an IGBT, which is industry standard for off-line these days. Unfortunately, although these transistors are driven like a MOSFET, you are then back to bipolar switching speeds again.

Continuous and Discontinuous

Continuous or discontinuous refers to the current in the inductor: in a discontinuous mode converter, the inductor current goes to zero at some time during the period. Stated differently, the difference between continuous and discontinuous mode is that to have continuous mode, you have to have enough inductance that at minimum load (including any preloading) there is still inductor current flowing at all times. In equations:

$$I_{\text{load,min}} \geqslant \frac{V_{\text{out}} T(1-D)}{L}$$

where T is the period and D the duty cycle, and we have assumed that the forward voltage drop of the rectifier is small compared with the output voltage. Of course, if minimum load current is zero, you necessarily have discontinuous mode (except see below).

> **Practical Note** The key thing is to choose <u>either</u> continuous or discontinuous; don't allow the converter to be sometimes one and sometimes the other depending on the load. This can make it difficult to stabilize the loop.

An exception to this general rule occurs with synchronous rectification. A converter using synchronous rectification is <u>always</u> in continuous mode. Thus, no minimum inductance is required.

Synchronous Rectification

In many applications nowadays, converter efficiency is (almost!) more important than cost. Indeed, looked at from the consumer's viewpoint, a more efficient but more expensive up-front converter actually is cheaper, because the cost of downtime can be so high: an extra half-hour of compute time on a laptop computer, for example, would certainly be worth an extra dollar in the power supply.

When efficiency is important, it certainly pays to consider the use of a synchronous rectifier, that is, a system in which the function of the output rectifier is accomplished by a switch, invariably a MOSFET. Many ICs available today will drive both a main switching FET and a synchronous rectifier as well, so this can be far less painful than it was just a few years ago, when a second drive had to be developed using discrete components.

A further reason to consider using synchronous rectification is that as noted above, it converts a (potentially) discontinuous mode operation converter into a continuous mode operation. This is because even at no load, the current can flow in either direction in the inductor (because an "on" MOSFET can conduct in both directions). Using a synchronous rectifier, then, relieves you of having to worry about changing modes (which can be bad for converter stability, see Chapter 6), or about minimum inductance to ensure continuous operation.

One small downside to synchronous rectification deserves mention here. The main switching MOSFET has to be off before the synchronous rectifier is turned on, and vice versa. If this detail is neglected, there will be shoot-through: the input (or output) voltage will have a direct path to ground, engendering very high losses and potential failure. During the interval of time when both MOSFETs are off, the current in the inductor has to flow *somewhere*. Generally, the body diode of the MOSFET should not be used to carry this current, because this diode has a very long reverse recovery time. Suppose the body diode is intended carry the current while the MOSFETs are off. While the body diode is recovering, it acts like a short, so there is once again a path from input (or output) to ground, giving rise to shoot-through. To get an idea of the potential for trouble here, consider Figure 2.1*B*.

The bottom line is that it is necessary to have a schottky diode in parallel with the MOSFET's body diode, to carry the current during the time when both FETs are off. (The schottky has a much lower V_f than the body diode, and so carries essentially all the current; the reverse recovery time of the body diode depends on its previous forward conduction current, which is therefore negligible.)

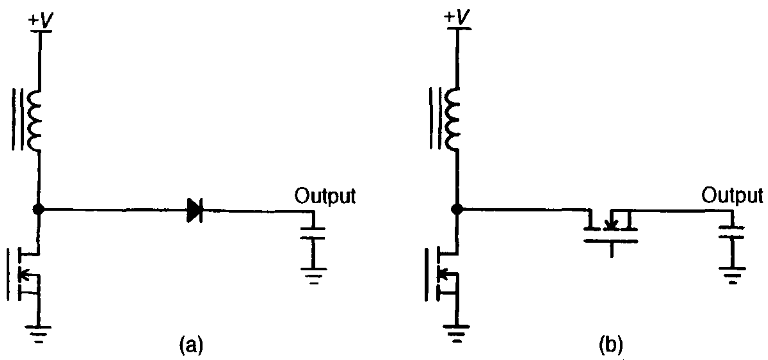

Figure 2.1 (a) Nonsynchronous converter uses a diode, whereas (b) a synchronous converter uses a MOSFET.

Voltage Mode versus Current Mode

You might observe that the distinction between current mode control and voltage mode control has not been mentioned in this list of things to consider up front. This is because this is really a control issue; most every topology can have either type of control. There is one point, though, on which there can be significant effects of selecting one or the other: if currents are high, current mode is going to require sensing the current with either a resistor (which will dissipate a lot of power) or a current transformer (which costs money). As a mitigating factor, though, this sensing makes overcurrent limiting straightforward. So for higher power outputs, it's worth thinking about this choice as well.

Conclusions

The more you know about the system in which your power supply is going to be operating, the better you can make design choices up front. And making proper choices at the beginning is vastly less costly and time-consuming than trying to make fixes later.

> **Practical Note** Make yourself a checklist from the specification sheet for the converter, and go through each of the items above. You'll often find that you come down to only one or two possibilites for a topology based on these constraints, and then topology selection may be easy, based on cost or size. For convenience, Table 2.1 summarizes the various choices talked about in this section.

TABLE 2.1 Topology Selection Checklist

1. *Step-up or step-down.* Is the input voltage always higher or always lower than the output? If not, you can't use a buck or non-isolated flyback.
2. *Duty cycle.* Is the output voltage different by more than a factor of 5 from the input voltage? If so, you probably will need a transformer. Calculate duty cycle to ensure that it doesn't have to get too small or too large.
3. *How many output voltages are required?* If more than one, a transformer may be required, unless you can postregulate. Large numbers of outputs suggest more than one converter may be a good choice.
4. *Is isolation required?* How much voltage? Isolation necessitates a transformer.
5. *What are the EMI requirements?* Tight requirements suggest staying away from topologies with discontinuous input current, such as a buck, and choosing continuous mode operation.
6. *Is cost so paramount that a BJT might be a choice?* Or if off-line, an IGBT? Otherwise, plan on MOSFETs.
7. *Is the supply required to operate with no load?* If so, choose discontinuous mode—unless the answer to question 8 is yes:
8. *Can synchronous rectification be afforded?* This makes the converter continuous regardless of load.
9. *Is the output current very high?* Then it might be good to use voltage mode rather than current mode.

THE BUCK TOPOLOGY

Turning from generalities to specific converters now, it is assumed that you know what a buck converter looks like. A sample is shown later (Figure 6.17). Instead of being yet another compendium of topologies, this section, and those following it on the other topologies, concentrates on practical difficulties with each topology, and some possibilities for circumventing them. Concentrating on the problems up front will enable you to make a better selection of topology, by highlighting areas that will consume much of your time in the design and debugging phases.

Limitations

As mentioned under General Considerations, there are a number of limitations to the buck topology that need to be addressed at the start.

1. Although a buck converter is conceptually clean in having only an inductor and no transformer, this means in turn that it's not possible to have input-to-output isolation.
2. The buck can only step down the input voltage: if the input is ever less than the desired output, the converter won't work. (However, see the section below on the buck-boost.) You can use a buck to generate a negative voltage. Figure 2.2 shows such a configuration. When the transistor turns on, current in the inductor ramps up. When the transistor turns off, the inductor current is pulled from the output capacitor, pulling it negative.

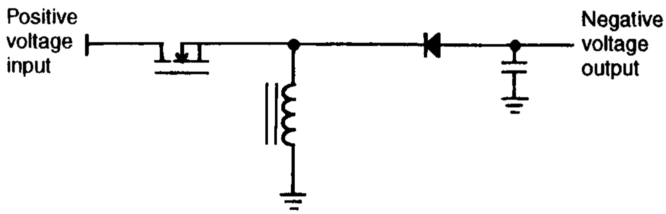

Figure 2.2 Using a buck to convert a positive input voltage to a negative output voltage.

3. The buck only has one output. This is fine if you're looking for a 5V-to-3.3V converter, but unless you're willing to contemplate a second stage of regulation, such as a linear postregulator, the many applications in which you're looking for multiple outputs are ruled out.
4. Although the buck can be either continuous or discontinuous, its input current is always discontinuous, meaning that during the portion of the cycle when the transistor is off, the input current goes to zero. This makes the EMI filter larger than it might need to be with other topologies.

Gate Drive Difficulties

Driving the gate of a buck can get to be quite a nuisance, not to say a problem. The trouble is that to turn on an n-channel MOSFET, the gate voltage has to be at least 5V and more likely 10V above the input voltage (respectively 1V and 5V for logic-level FETs). But how do you generate a voltage higher than the input? The easiest way around this problem is no doubt to use a p-channel FET, so it can be turned on just by pulling the gate to ground. Unfortunately, p-channel FETs usually have substantially higher $R_{DS,on}$ than n-channels do, and cost rather more. Besides, the input voltage would have to be less than 20V to avoid blowing out the gate, ruling it out in a number of applications. The reality of using p-channel MOSFETs is this: with a pull-down resistor, you usually can't get enough switching speed on the gate for the efficiency you want, and you end up going back to an n-channel after a few frustrating days of lab work.

Practical Note Except for very low input voltage converters, build your buck converter with an n-channel MOSFET.

One common way to drive the gate is to use a gate drive transformer that isolates the driver from the gate (Figure 2.3).

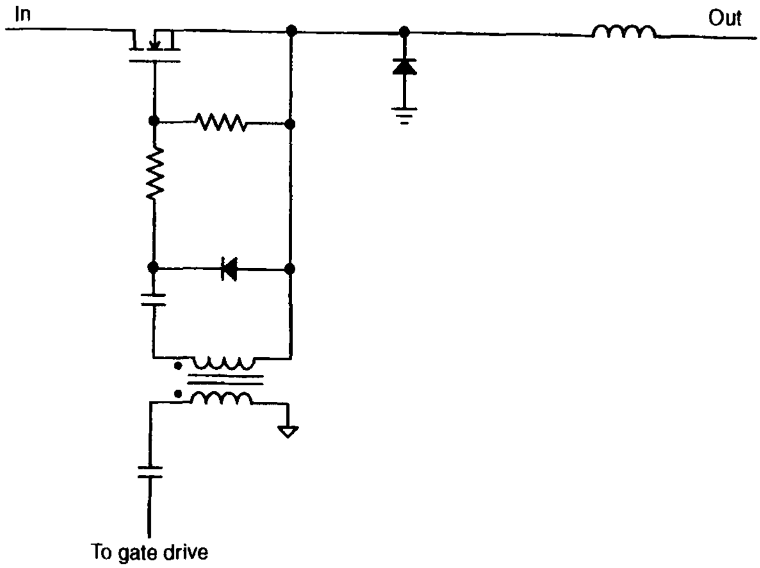

Figure 2.3 Use of a transformer to drive a buck transistor.

The capacitor on the drive side of the isolation transformer prevents DC current from flowing through the primary while the gate drive output is high. The capacitor and diode on the other side restore the voltage to unidirectionality—otherwise a 12V drive on the primary becomes a ±6V drive on the secondary. The gate resistor is always necessary (see the discussion in Chapter 3 on components), and finally, the gate–source resistor is just a bleed: if the gate drive stops switching for some reason, the gate eventually turns off.

Practical Note Choose the two capacitors in this gate drive circuit to be at least 10 times bigger than the gate capacitance—remember that they form a divider with this capacitance, and so this way you'll get at least 90% of the drive voltage on the gate.

Although this system is relatively cheap and works well, it is limited in maximum duty cycle because the transformer has to have time to reset.

A method that allows extremely fast gate drives utilizes a separate push–pull housekeeping converter to generate a DC secondary voltage referenced to the source of the MOSFET (*F* ground in Figure 2.4). It's not necessary for this second converter to be in a closed loop; if it comes from a regulated source, a fixed duty cycle converter works well. You can then have a gate driver IC referenced to the source, and really drive the MOSFET fast. Although I have used this circuit many times, it is somewhat expensive because of all the extra parts needed. (You could use a 555 timer for generating the 50% duty cycle.)

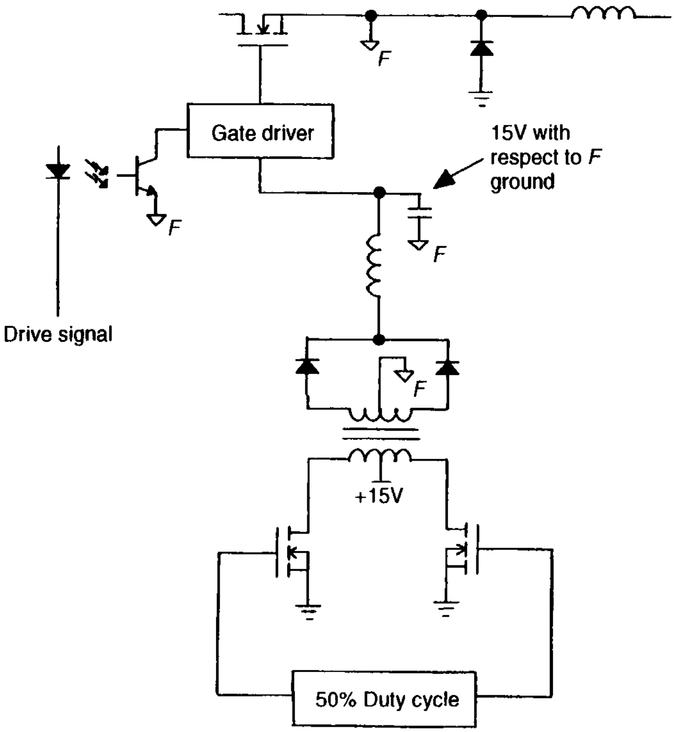

Figure 2.4 Generating a floating supply to drive a buck transistor.

You also need a way of signaling the floating system to control the gate driver. The signal of course can't tolerate excessive propagation delay, ruling out slow optocouplers such as the 4N48. To avoid having yet another transformer, I have found the HCPL2601 family of optocouplers to be excellent even for very high input voltages, because of their excellent dV/dt rating.

THE FLYBACK

Two Kinds

There are two kinds of flyback, the nonisolated flyback (Figure 2.5) and the isolated flyback (Figure 2.6),

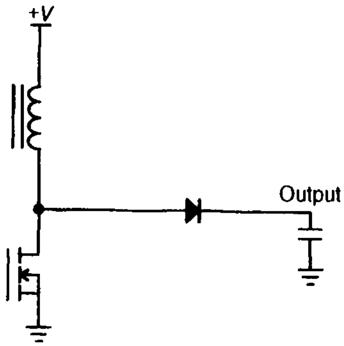

Figure 2.5 Basic nonisolated flyback topology.

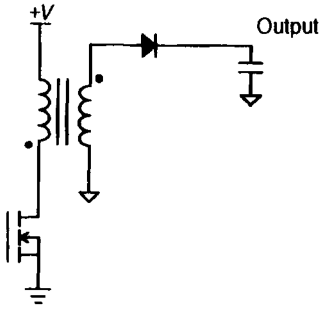

Figure 2.6 Basic isolated flyback topology.

which we explicitly show to avoid name confusion (see below). To be absolutely sure, let's briefly describe their operation.

The nonisolated flyback turns on its switch for a fraction D of the switching period, which, since it produces a voltage across the inductor, causes current to ramp up, storing energy in the inductance. (More explicit details are given in Chapter 5 on magnetics design.) When the switch turns off, the inductor current goes through the diode and into the output capacitor and load.

The isolated flyback works entirely analogously. During the on-time of the switch, energy is stored in the inductance of the primary. Looking at the dots on the transformer, we see that when the switch turns off, the drain voltage rises above the input voltage, which causes the secondary voltage to rise above ground; this turns on the diode, again providing output current to the capacitor and load.

The nonisolated flyback has a single output (there's no way to make more than one). That output is not isolated from the input, and the output can't be made less than the input voltage—even if you turn the transistor completely off, the output will equal the input (minus a diode drop). On the other hand, if all you're looking for is a single nonisolated output, this flyback has only a single-winding inductor to deal with.

The isolated flyback can have multiple outputs if multiple secondaries are put on the transformer, and all those outputs can be isolated from the primary, and potentially from each other. Further, the outputs can be made to have any value whatsoever, simply by adjusting the primary to secondary turns ratio. The downside is that the magnetics is now a multiwinding transformer (see below).

Name Confusion with Boost

Frequently people call the nonisolated flyback a "boost" converter. The term "boost" does not appear again in this book. The terms "nonisolated flyback" and "isolated flyback" are consistently used to refer to the topologies shown in Figures 2.5 and 2.6. As discussed in Chapter 6, the distinguishing feature of a flyback topology is that the magnetic structure stores energy during a portion of the switching cycle; this is why we use the same name for these two topologies.

Practical Note Whenever you're reading something that refers to a "flyback" or "boost," look carefully at the schematic to see what topology is actually being talked about. The literature is inconsistent, resulting in endless confusion.

Continuous versus Discontinuous

Both types of flyback can be run in either continuous or discontinuous mode. In a general sort of way, though, a flyback is usually used to enable the converter to go to no-load current without needing any preload. (At no load, the switch is simply turned off until the charge on the output capacitor bleeds off, and then turns on again for a single pulse. This is known variously as "pulse-skipping mode" and other similar terms.) For this no-load operation to work, you need to operate in discontinuous mode, and as indicated before, it's best not to change modes because of difficulties in controlling the converter's loop. The most common operation of a flyback is thus in discontinuous mode.

Capacitor Limitations

When the flyback transistor turns off (see the discussion in the magnetics chapter, Chapter 5, for more on this), the energy stored in the primary inductance comes out on the secondary winding(s). Since there is no inductor on the secondary, the full peak current goes straight into the capacitor. At higher power levels, it can become quite hard to find a capacitor with sufficient ripple current rating to handle this: remember that you have to calculate the RMS current to know whether the capacitor can handle it. Suppose for example that we are running a 5V output at 10A (this is about the limit for a flyback, see below), and the duty cycle is 50% at this power level. The transformer has to deliver the 50W for the full period in just half the period (since the duty cycle is 50%), so the current it delivers during the conduction time of the diode is double, 20A. So the RMS current is

$$I_{\text{RMS}} = \sqrt{\tfrac{1}{2}(20\text{A})^2} = 14\text{A}$$

This extremely high current will require paralleling many aluminum or tantalum capacitors, or else the use of a high-priced MLC cap. Failure to get adequate capacitance on the output of a flyback is a major cause of capacitor failure.

Power Limits

There is a maximum power that can usefully be used with a flyback, on the order of 50W for low voltage inputs. (You sometimes hear stories from people who say they built one at 500W, but they don't tell you that it could never be made to work on the production line.) In any event the power output is inversely proportional to the inductance; to get a large power requires a tiny inductance (the math is detailed in the chapter on magnetics). By the time you get up to 50W at a reasonable switching frequency, the inductance is very small (the same order of magnitude as strays); this makes the design almost impossible to produce consistently in production. For example, a slight change in the lay of the wires by the magnetics vendor will affect the inductance enough to prevent you from getting maximum power out.

Tip For low voltage inputs, limit flybacks to designs requiring less than 50W, somewhat more for high voltage inputs.

Practical Limits on Number of Outputs

Of course, for all converters, the transformer becomes more difficult to wind as you add more windings to it. For an isolated flyback, however, this difficulty is crucial. The regulation of each output depends on the leakage inductance of the winding, because the leakage inductance subtracts from the voltage delivered to the output. So to get good tolerance on the outputs, the leakages must be either negligibly small (almost impossible) or the same from unit to unit, so that they can be compensated for. If you have multiple windings, however, controlling (and even measuring and specifying) the leakage on each winding is almost impossible. The author once saw an isolated flyback design with (believe it or not!) 13 outputs. According to the designer, a flyback "was cheaper than a forward because it didn't need an inductor." Unfortunately, after this was in production, the vendor's winding person (there was only one vendor, no one else wanted to touch it) left the magnetics company, and thereafter no one else was ever able to wind the transformer in a way that made the circuit work!

Practical Note If you need more than three or four outputs, don't use a flyback. It will be cheaper in the long run to go to a forward.

THE BUCK-BOOST

"Buck-boost" is the standard name for what might be better called, in line with the terminology used in this book, a "buck-flyback." It's not a very common topology yet, but it has some advantages that suggest it will become increasingly used.

The buck-boost converter, as its name suggests, works as either a buck or a flyback, depending on whether the input voltage is, respectively, higher or lower than the output voltage. The great thing about this topology (Figure 2.7) is that this transition is accomplished automatically, there are no discrete changes involved.

In the buck-boost, both switches are on at the same time, and both are off at the same time. Consider first the case of the input voltage higher than the output. The top transistor

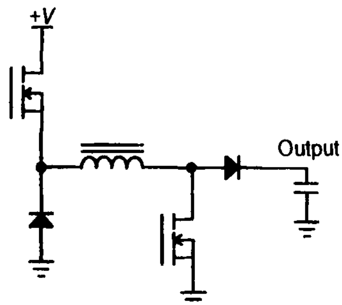

Figure 2.7 A nonisolated buck-boost.

(see Figure 2.7) acts as a buck switch, with the grounded-anode diode as the freewheeling rectifier. Since the bottom switch is on at the same time as the top switch, the full input voltage is applied across the inductor, ramping up the current. When both switches turn off, the grounded-anode diode carries the current, and the other diode simply forward conducts. This is thus a buck converter.

Next, suppose the input voltage is lower than the output. The ground-referenced transistor now acts as a flyback switch, and the second diode acts as the freewheeling rectifier. Once again, since both switches are on at the same time, the full voltage is applied across the inductor during the on-time.

Observe what the description has said: in both cases, whether acting as a buck or acting as a flyback, the full input voltage is applied across the inductor. But this means that *the same control circuit works for both "modes"* and consequently the converter does *not* switch between modes; therefore, stabilization of the loop is straightforward!

Limitations of the Buck-Boost

As we would expect, the problems with the buck-boost are a combination of the problems with the buck and the flyback. Acting as a buck converter, it has no input–output isolation, and there is only a single output. Acting as a (nonisolated) flyback, there's a maximum practical output power. And finally, unless you can replace the (schottky) diodes with two more MOSFETs to make the converter synchronous, there can be relatively poor efficiency; but a driver with four outputs (perhaps a full-bridge PWM IC?) would be required to achieve synchronous rectification. Still, the ability to work over a large range of input voltages, and the appearance of ICs for controlling this topology, may make the buck-boost an attractive choice.

THE FORWARD

Again, to avoid confusion with the term "boost," when this book refers to a forward converter, the topology illustrated in Figure 2.8 is always meant.

The forward works entirely differently from the similar appearing flyback. The key point is noticing that the dots on the transformer now mean that the output diode is forward-biased when the voltage across the primary is positive, that is, when the transistor is *on*; a flyback's diode is on when the switching transistor is off. Energy is thus not (intentionally) stored in the primary inductance, as it was for the flyback; the transformer

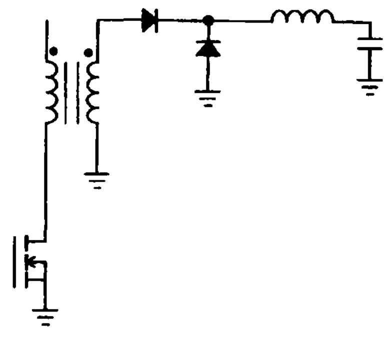

Figure 2.8 Basic forward topology.

acts strictly as a transformer. When the transistor is turned off, the only energy stored is that in the leakage inductance of the transformer; this is what causes the drain voltage to rise above the input voltage, resetting the core.

Minimum Load

The forward is one of those converters mentioned at the beginning of the chapter that requires a minimum load. The inductor has to be big enough to ensure that its peak ripple current is less than the minimum load current. Otherwise it will go discontinuous, and the output voltage will rise, peak detecting. This means that a forward converter cannot operate with no load, because you cannot have an infinite inductance.

> **Practical Note** A swinging choke, such as that produced with an MPP core, is an excellent choice here. A swinging choke is one whose inductance decreases gradually as the current through it increases. At minimum load, you get a lot of inductance, keeping the core continuous, and at maximum load you still have some inductance, but not as much; you allow the ripple to increase as the load current increases, so that the inductor doesn't have to be designed as physically big as would be needed to maintain the full inductance at maximum load.

One commonly used way around a minimum load is to attach a small load resistor (a "preload") permanently at the output terminals, as a part of the converter itself. Then, even when the external load is zero, the converter can remain continuous because it is still supplying some minimum power to this resistor. Of course this eats up a certain amount of power when the external load is above minimum.

> **Practical Note** Schemes abound for turning off this preload as external load increases. Very frequently, the result is oscillations: the preload turns off, then the converter goes discontinuous, which causes the preload to turn on, and the converter is continuous, causing the preload to turn off, etc. Just bite the bullet and accept the small efficiency hit compared with the cost (and efficiency hit) of a much larger inductor.

Leakage Inductance

Unlike the flyback, which uses its primary inductance to store energy, the forward really has parasitic leakage inductance. When current is flowing through the primary, there is energy stored in the leakage inductance, $\frac{1}{2}L_{\text{leakage}}\,I^2$. This energy has to go someplace. In the simplest case, you just throw it away, either into an *RC* snubber, or into the transistor itself, letting it avalanche. More sophisticated schemes recover more or less of the energy, using an additional winding on the transformer (though this doesn't work perfectly either because of leakage!) or some form of switched reactance, often using another FET. Regardless of what is done with the energy, it is a nuisance and to some degree an efficiency hit; the best approach is to wind the magnetics in a way that minimizes the leakage inductance.

Summary

Because the forward doesn't store energy in the transformer, it doesn't have the limitation that hinders the flyback in terms of power level; it also has an inductor, which smoothes the current seen by the output capacitors. Forwards can be straightforwardly constructed at a level of 500W or more. The main limitation of the topology eventually comes about, rather, from the available size of MOSFETs. Increased power translates into increased currents, and eventually losses in the MOSFETs become unacceptable. When this is the case, a topology with more than one transistor to share the burden is desirable.

THE PUSH–PULL

There are two basic styles of push–pull converter, current-fed and voltage-fed. The difference between them boils down to much nicer waveforms and operations in the current-fed push–pull, but at the price of having a (sometimes rather large) extra inductor.

The push–pull is treated here, while the half-bridge isn't, because the push–pull has both its transistors ground-referenced. Although it was noted above that there are ICs available that will drive high-side transistors for synchronous rectifiers, they tend to have rather low maximum voltages. Since the push–pull and half-bridge use two transistors, presumably they have been selected because the power level is higher than in single-transistor topologies, which often means that the input voltage is higher. Driving a half-bridge may thus get back into discrete parts to generate the floating gate drive; the push–pull definitely has an advantage here.

Voltage-Fed

The voltage-fed push–pull converter works by having two transistors across a center-tapped transformer (see Figure 2.9). They are operated 180° out of phase. This *doesn't* mean that each one is on 50% of the time, just that they have the same duty cycle, with one going on half a switching period later than the other. If the left transistor in Figure 2.9 is on, the right transistor is off. Looking at the dots on the transformer, this means that V_{in} is applied across half the transformer, and so $2 \times V_{in}$ is applied on the drain of the off transistor. Continuing with the left transistor on, there is a positive voltage applied to the bottom diode, which is on, and the top diode is reverse-biased. Everything is then mirror-

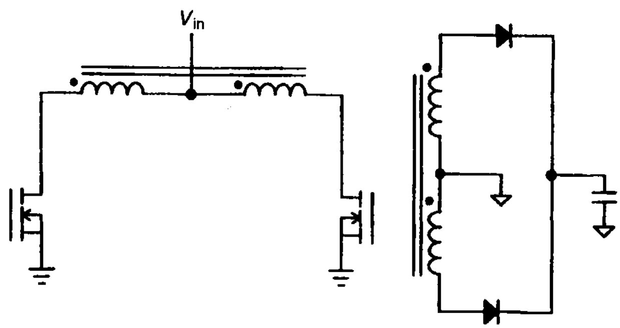

Figure 2.9 A voltage-fed push–pull.

imaged when the right transistor is on; since the two transistors are on for the same amount of time, if V_{in} is constant during a switching period, the volt-seconds across the transformer ideally sum up to zero, and the core operates symmetrically around zero gauss.

The biggest problem with this converter is the voltage rating needed by the transistors, which is at least double the maximum V_{in}. Operating from a rectified 120V line means that the transistors will see at least $2 \times 120\text{V} = 240\text{V}$. In practice, the line can be a very nasty place, as noted in Chapter 1, and so a 400V transistor might be a common choice here. This high voltage rating in turn means that the $R_{DS,on}$ is high, and so losses may be higher than desired. And in any case, the V_{in} need only surge above 200V for one switching period to blow out the transistors!

The other potential problem is that there must be a time (the dead time) between turning off one transistor and turning on the other. If both transistors were on at the same time, the transformer would be effectively shorted, and so the current would rapidly increase, limited only by the leakage inductance—this is a common cause of failure. The transistors must also be on equal amounts of time to avoid saturating the transformer; in practice this is accomplished by using current mode control (see Chapter 5 for the concept of saturation, and Chapter 6 for current mode control).

Current-Fed

The sensitivity to line voltage exhibited by the voltage-fed push–pull is obviated in the current-fed push–pull because it has an inductor between V_{in} and the transformer. Now when the transistor turns on, it gets a current set by the inductor current, as shown in Figure 2.10. This arrangement also gets rid of the problem of having to turn off one transistor before the other turns on, since even if both transistors are on simultaneously, the current is still limited by the inductor.

The downside of this converter is the addition of an extra inductor. Since this device must both carry the DC current of the converter and provide sufficient inductance to act like a current source during a switching period, it can easily grow to rather large (read expensive) size for moderate power level converters.

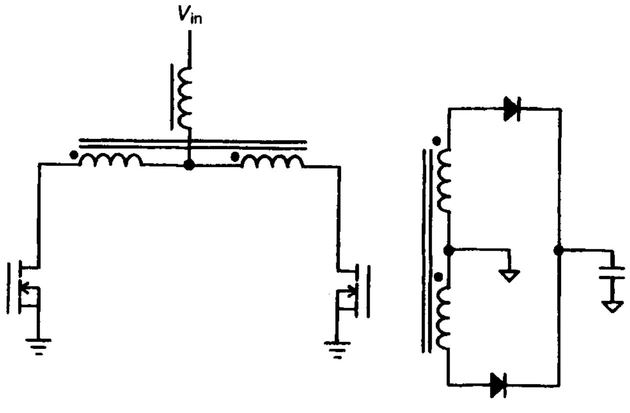

Figure 2.10 A current-fed push–pull.

Transformer Utilization

It should be observed that all the topologies discussed up till this section (the flyback, the forward, and the buck-boost) utilize only half the magnetics' cores: the flux density is ramped up to a maximum value and then back down to zero, never going negative. The push–pull utilizes the magnetics better, because the core's flux density goes both positive and negative, thus reducing the size of the magnetics for a given power level compared to the single transistor topologies.

RESONANT CONVERTERS AND SOFT-SWITCHING CONVERTERS

For quite a few years now, everywhere you turn there have been articles about resonant converters, how great they are, and how everyone ought to be using them. (The author feels, however, that this fad is finally passing.) If you are one of the braver souls, perhaps you've actually been in the lab, and spent several weeks or months trying to make a go of a resonant converter.

By way of contrast, there seems to be very little heard about soft-switching converters, and yet they seem much more practical; many of the converters in production that are called resonant are actually soft-switching. Another name for soft-switching is "quasi-resonant."

As noted at the start of this chapter, there are hundreds of different topologies that are resonant or soft-switching; for this reason, this section merely points to the sorts of feature it might benefit you to investigate.

The Difference Between Resonant and Soft-Switching Converters

A resonant converter is one in which the power waveforms (current and voltage) are sinusoidal. This is accomplished by letting inductances form a resonant tank with capacitances, the latter often (though not always) being parasitics. Switching occurs when the voltage and/or current goes through zero, ensuring an almost lossless switch transition. Resonant converters thus have had their main claim for usefulness in high frequency converters, where switching losses can dominate on-state losses of the switches. However, since switch transitions depend on the frequencies of resonant tanks, the actual switching frequency of the converter varies, sometimes quite dramatically, usually as a function of load and line.

A soft-switching converter is intermediate between a resonant converter and a PWM converter. Any of the topologies described in the sections above can be made soft-switching by suitable addition of components. A soft-switching converter always switches at the same constant frequency, like a PWM, but it creates a tank circuit for a portion of the switching period so that switch transitions still occur nearly losslessly.

Why You Should Not Use Resonant Converters

Resonant converters have quite a number of problems. Not least among these is the variation of switching frequency with load. In fact, for a common class of these converters, minimum switching frequency occurs at maximum load, so that EMI filtering has to be designed for the worst combination, minimum frequency and maximum current. The gain in size from operating at a high switching frequency may be lost when a realistic converter is designed including the EMI filter. The next time you are told about a resonant converter that does $100W/in^3$, ask what the power density is when a noise filter is included!

An even more serious problem arises because of the common use of capacitive strays as one of the elements of the resonant tank. This strategy almost *can't* be made to work on a production line, although it's great in the lab. The trouble is that these strays are not consistent from device to device; they can even differ between two *identical* devices from different manufacturers! This variation directly affects the operational frequency, which affects the output caps, the EMI filter, etc. The only way around it is to parallel some external capacitance with the parasitic, so variation of the parasitic is relatively unimportant. Unfortunately, this modification increases the tank period, and so the original motivation, operating at a high frequency, is destroyed.

Why You Should Use Soft-Switching Converters

In contrast with resonant converters, soft-switching converters operate at a fixed frequency, making their filtering requirements straightforward. They also typically use discrete capacitors, and so have quite reproducible characteristics from unit to unit. Figure 2.11 shows a fairly standard implementation of a soft-switching forward converter, with a sketch of a drain waveform.

Initially, the transistor is on, and the drain voltage is zero. When the transistor turns off, the primary inductance of the transformer forms a resonant tank with the external capacitor (in parallel with the drain–source capacitance of the MOSFET, but the external capacitor is designed to be larger than the MOSFET's capacitor). After completing a half-cycle of the ring, the core is reset: the L and C values set the ring frequency, and the volt-seconds required to reset the core determine then how high the voltage rings up. After the half-cycle ring is completed, the drain voltage remains at the input voltage, since there is now no energy stored in the transformer. It remains in this state until the transistor turns on again.

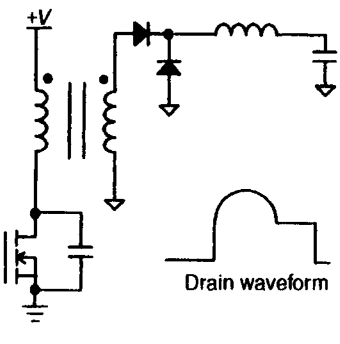

Figure 2.11 A quasi-resonant or soft-switching forward converter.

What distinguishes this converter from a resonant converter is that it is still pulse width modulated: the transistor has a constant switching frequency. Of course, the capacitance and the inductance still have to be chosen carefully. If they are too big, the (half) period will exceed the switching period, and the core won't reset; if they are too small, the drain voltage will go excessively high, to get the necessary volt-seconds in a very short time. Even so, there is wide room for variation of the stray components within which the converter will work normally.

It may be noted that when the transistor turns on, the capacitor energy is dissipated into the MOSFET. If the capacitor is sufficiently small, however, this may not be too terrible. For example, if the capacitor is 100pF, the input voltage is 50V, and the switching frequency is 500kHz, the power lost because of the capacitor is only $P = (\frac{1}{2}) \times 100\text{pF} \times (50\text{V})^2 \times 500\text{kHz} = 63\text{mW}$.

Indeed, the only bad thing about soft-switching converters is the apparent dearth of ICs designed to operate them, although something can be rigged from a PWM designed to operate a synchronous rectifier. Perhaps as the word gets out (and certain potential patent issues become clarified) ICs implementing soft-switching will become common—at that time, they will make an excellent choice.

COMPOUND CONVERTERS

A compound converter is any converter that has two (or in theory more) stages in series. It is distinguished from merely two series converters in that there is usually only one control loop for the whole system. For example, one possible compound converter consists of a front-end buck operating from 160VDC followed by a push–pull (see Figure 2.12). The buck operates closed loop to produce an approximately fixed output voltage (say 50V); the push–pull operates at a fixed duty cycle to step down the voltage (say to 5V). The loop is closed by sensing the 5V output, and using its error signal to control the duty cycle of the buck. Thus, although the push–pull is seemingly operated open loop (since it switches at a

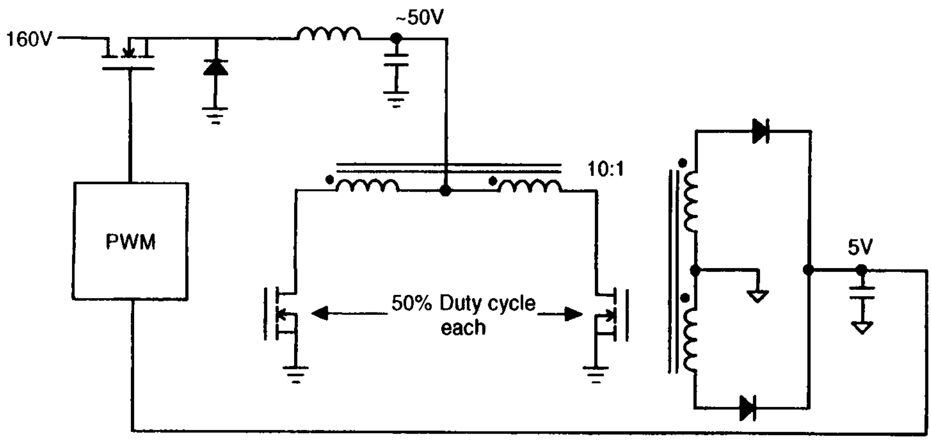

Figure 2.12 A compound converter consisting of a buck followed by a voltage-fed push–pull; the buck's output capacitor serves as the push–pull's input capacitor.

fixed duty cycle) it is actually just a gain block inside the control loop, which is closed around it (in the example shown in Figure 2.12, it has a gain of $1/10 = -20$dB).

In some cases, the two converter stages may share components; in the example just given, the output capacitor of the buck is also used as the input capacitor of the push–pull. It is easy to imagine other combinations in which an inductor could be shared instead. As with the resonant and soft-switching converters, there are a large number of possible combinations for compound converters; instead of attempting an enumeration, I'll comment about when they might be useful.

When to Use Them

As the example given shows, having a compound converter is useful when you want to get a lot of step-down or step-up. It's already been mentioned that there are practical limits to the duty cycle you can get from a PWM, and to the size of the turns ratio you should try for on a transformer. If you need to make a voltage conversion beyond what's feasible within these limits, a compound converter offers a way to considerably extend the transformation range available.

A compound converter might be desirable, as well, when you need to get a fairly large conversion ratio (of input to output voltage) in a situation that also calls for input-to-output isolation. The two requirements together can make for a very challenging design, but by segregating the functions, you can make it much easier: For example, let the front-end converter do the voltage transformation, and then let the second converter do the isolation, perhaps with a 1:1 transformer. Since the second converter would always operate with the same input voltage and the same output voltage, its components could be optimized for this operation, and it could be very efficient. Indeed, this compound convertor may well be more efficient than a single-stage converter, because of the difficulties in the magnetics involved with simultaneous design of both large conversion ratio and isolation.

REFERENCES

1. Issa Batarseh, *IEEE Transactions on Power Electronics*, PE-9(1), 6 (1994).