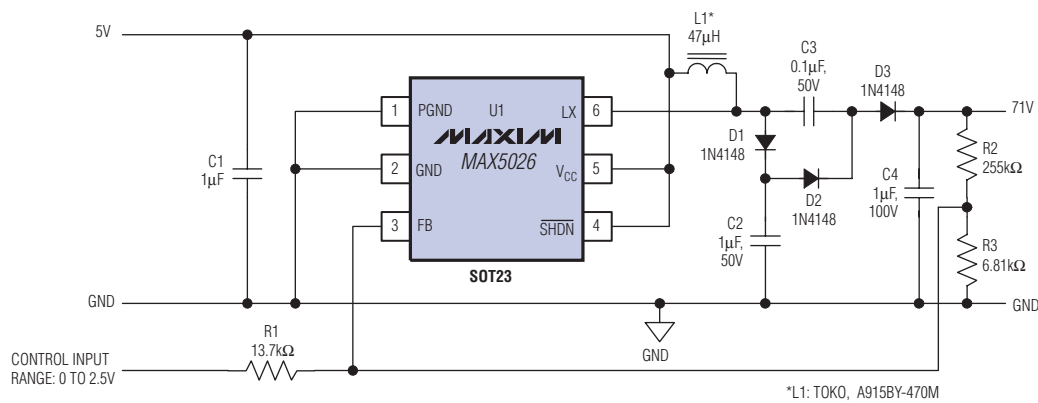


|                          |   |           |
|--------------------------|---|-----------|
| <b>NEWS BRIEFS</b>       |   | <b>2</b>  |
| <b>IN-DEPTH ARTICLES</b> | <b>Low-noise APD bias circuit</b>   | <b>3</b>  |
|                          | <b>Understanding successive approximation ADCs</b>                                    | <b>6</b>  |
|                          | <b>Calculating MOSFET power dissipation for high-power, portable DC-DC converters</b> | <b>10</b> |
| <b>DESIGN SHOWCASE</b>   | <b>Power meter achieves <math>\pm 1\%</math> accuracy</b>                             | <b>16</b> |
|                          | <b>Portable devices derive 3.3V and 5V power supplies from USB ports</b>              | <b>18</b> |

## LOW-NOISE APD BIAS CIRCUIT



A control range of 0 to 2.5V adjusts the bias circuit output between 24.7V and 71V for desired APD gain.

---

---

# News Briefs

---

---

## **MAXIM REPORTS REVENUES AND EARNINGS FOR THE FOURTH QUARTER AND FISCAL YEAR**

Maxim Integrated Products, Inc., (Nasdaq: MXIM) reported net revenues of \$280.1 million for its fiscal fourth quarter ending June 29, 2002, an 8.4% increase over the \$258.5 million reported for the third quarter of fiscal 2002. Diluted earnings per share were \$0.20 for the fourth quarter on net income of \$68.6 million compared to a loss of \$0.05 per share on a net loss of \$16.2 million reported for the fourth quarter of fiscal 2001 and diluted earnings per share of \$0.19 on net income of \$66.7 million reported for the third quarter of fiscal 2002.

For the fiscal year, Maxim reported net revenues of \$1.025 billion compared to \$1.577 billion for last year, a 35.0% decrease. Net income for the year was \$259.2 million with diluted earnings per share of \$0.73 compared to \$334.9 million and \$0.93 per share in fiscal 2001. Included in fiscal year 2001 fourth quarter results was a pretax charge of \$163.4 million for merger and special charges related to the acquisition of Dallas Semiconductor.

During the fourth quarter of fiscal 2002, the Company repurchased 7.3 million shares of its common stock for \$367.7 million and acquired \$23.0 million of capital equipment. Year-end cash, cash equivalents, and short-term investments were \$765.5 million. Accounts receivable increased by \$21.2 million in the fourth quarter to \$129.8 million, and inventories decreased \$8.5 million to \$139.2 million.

Gross margin decreased from 70.2% in the third quarter to 68.1% in the fourth quarter, primarily as a result of revenue growth in lower margin products and the recording of \$3.9 million of inventory reserves compared to \$2.1 million last quarter. Research and development expense increased to \$72.0 million or 25.7% of net revenues in the fourth quarter, compared to \$69.0 million or 26.7% of net revenues in the third quarter. The increase in research and development spending resulted from hiring additional engineers and increased spending to support new product development efforts. Selling, general, and administrative expenses remained relatively unchanged during the quarter.

Fourth-quarter bookings were approximately \$310 million, a 4% increase over the previous quarter's level of \$299 million. Turns orders received in the quarter were \$140 million (turns orders are customer orders that are for delivery within the same quarter and may result in revenue within the same quarter if the Company has available inventory that matches those orders). Bookings increased in Japan and the U.S. and were approximately flat with last quarter in other geographic areas. Fourth-quarter ending backlog shippable within the next 12 months was approximately \$239 million, including approximately \$210 million requested for shipment in the first quarter of fiscal 2003, which is up 7% from the previous quarter. The Company's third-quarter ending backlog shippable within the next 12 months was approximately \$219 million, including approximately \$195 million that was requested for shipment in the fourth quarter.

Jack Gifford, Chairman, President, and Chief Executive Officer, commented: "Fourth quarter results were generally consistent with our expectations. Revenues increased 8% and bookings increased over last quarter. Turns orders remained high at 45% of bookings. Our customers continue to order for their immediate needs. Although our backlog grew 9% in the fourth quarter, there is still limited visibility in many of our end markets.

"I am very proud of Maxim's accomplishments during what has been a very challenging year for most technology companies. We have grown our revenues for three consecutive quarters, and we forecast that revenues will be up sequentially for the first quarter of fiscal 2003. We believe that end-market consumption of our products is close to our current bookings level.

"We exceeded our original cost-savings expectations related to our acquisition of Dallas Semiconductor, achieving savings of over \$40 million per quarter over pre-acquisition spending levels. Approximately 75% of Dallas parts are now being tested at Maxim's Cavite facility in the Philippines, and we anticipate that over 80% will be tested there in fiscal 2003. Their gross margins have improved over 15 percentage points since the acquisition, and their new product launches and introductions more than doubled year over year."

Mr. Gifford concluded: "We hope, and we see some indications, that the business climate in the coming year is improving so that we can eventually increase our growth consistent with our past performance. Even if the business climate does not regain the full ebullience we experienced a couple of years ago, we at Maxim believe we can still deliver the outstanding profitability that our stockholders expect of us. We are committed to finding the right mix of operating efficiencies and market success that will produce those results for our stockholders."

# Low-noise APD bias circuit

Avalanche photodiodes (APDs) are used as receiving detectors in optical communications. The APD's high sensitivity and wide bandwidth make it popular with designers. APDs operate with a reverse voltage across the junction that enables the creation of electron-hole pairs in response to incident radiation. The electron-hole pairs are then swept by the applied field and converted to a current that is proportional to the radiation intensity.

Applying a variable reverse-bias voltage across the device junction creates a variable avalanche gain during APD operation. In turn, varying the avalanche gain optimizes sensitivity in the fiber-optic receiver. To achieve satisfactory levels of avalanche gain, however, many APDs require high reverse-bias voltages in the 40V to 60V range, and some require voltages as high as 80V.

A disadvantage of the APD is that avalanche gain depends on temperature and varies with the manufacturing process. Thus, for typical systems in which the APD must operate at constant gain, the high-voltage bias must vary to compensate for the effects of the temperature and manufacturing process on the avalanche gain. To achieve constant gain in a typical APD supply, the temperature coefficient must be maintained at approximately  $+0.2\%/^{\circ}\text{C}$ , which corresponds to  $100\text{mV}/^{\circ}\text{C}$ .

## APD power supply

Many methods exist for adjusting the output voltage of a power supply to compensate for temperature-induced gain variations of the APD. APD modules contain temperature-measuring devices such as thermistors, which can be connected directly to the power supply for output-voltage adjustment. In some systems, a microcontroller ( $\mu\text{C}$ ) reads the resistance value and then issues necessary bias-adjustment commands to the power supply.

The schematic of an APD bias power supply (**Figure 1**) is based on a low-noise, fixed-frequency PWM boost converter (U1) with an inductor that operates in discontinuous-current mode. Switching times have been intentionally slowed to reduce the high-frequency spikes that are otherwise present in most cases. Slower switching times reduce the high-frequency di/dt and dv/dt rates, which minimize radiated and coupled noise to surrounding circuits through current loops and capacitances between PC board traces or component pins.

Operating an inductor in discontinuous-current mode allows the decaying inductor current to naturally turn the diode to off. The MAX5026 switching frequency is 500kHz, and the internal, lateral-DMOS switching device has an absolute maximum rating of 40V. Operating with the external voltage-doubler network formed by C3, C4, D2, and D3, the circuit produces output voltages up to 71V.

Steady-state operation of the doubler circuit functions as follows: C2 transfers charge to C3 during the on time,

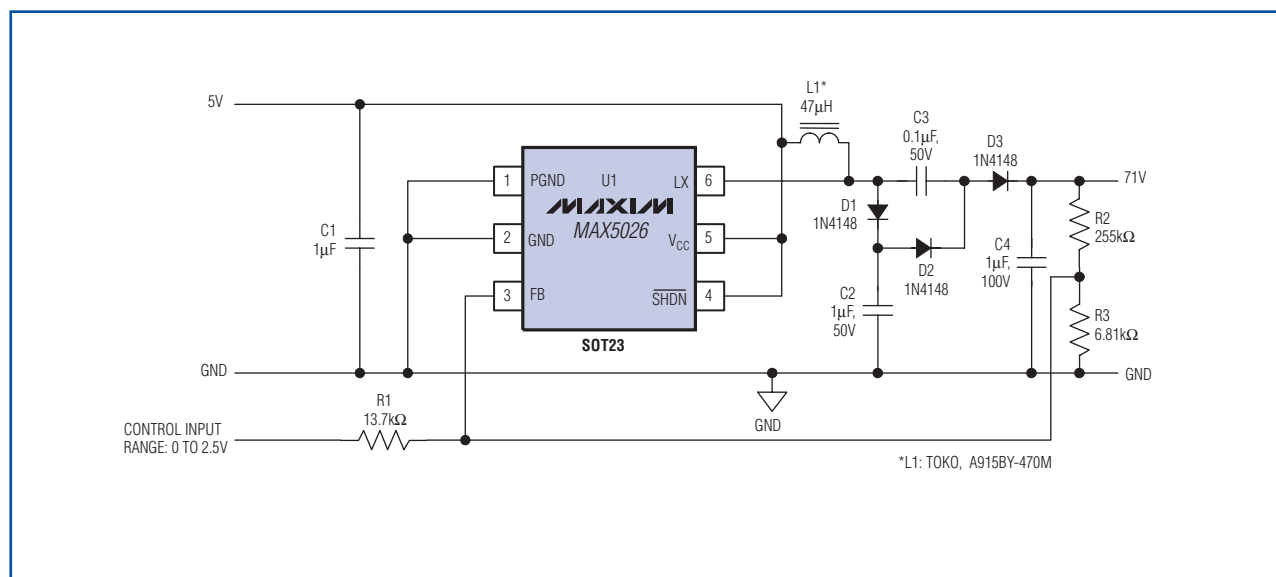


Figure 1. By varying the control input voltage from 0 to 2.5V, the low-noise APD bias power supply produces an output voltage change from 71V to 24.7V.

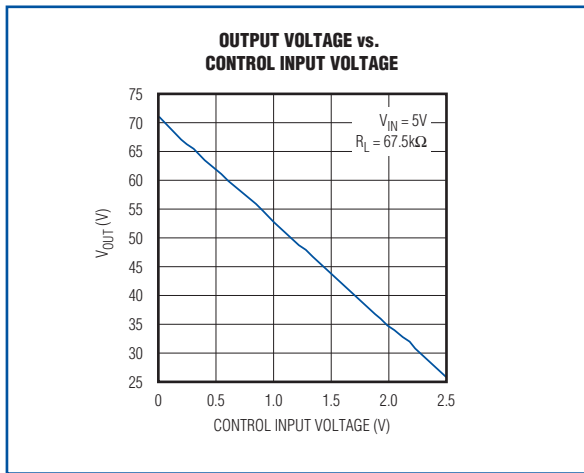


Figure 2. This graph demonstrates measured output voltage vs. control input voltage for Figure 1's circuit.

when L1 is charging and the LX pin is low (internal DMOS conducting). When the internal DMOS switches off, the inductor current forward biases D1 and D3. Thus, the total voltage presented to capacitor C4 is the sum of VC2 and VC3.

The MAX5026 benefits this application with:

- Slow rise and fall times at the internal FET minimize coupled di/dt and dv/dt noise.
- Discontinuous-mode inductor operation naturally commutates D1, virtually eliminating the high di/dt noise caused by reverse recovery in the diode.
- Fixed-frequency, 500kHz PWM operation generates a predictable and easily filtered noise spectrum.
- High integration results in low cost and small size.




Figure 3. This graph shows the efficiency curves vs. output current of Figure 1's circuit.

With a 5V input, Figure 1's circuit provides more than 1mA of output current at 71V output. **Figure 2** shows the range of output-voltage adjustment with respect to the control input voltage, and **Figure 3** shows efficiency curves for three output-voltage settings. Set the circuit's output voltage as follows:

$$V_{OUT} = \frac{V_{REF} \times (R2 \times R3 + R1 \times R2 + R1 \times R3) - V_C \times R2 \times R3}{R1 \times R3}$$

where  $V_{REF} = 1.25V$  and  $V_C$  is the control input voltage.

Figure 1's circuit has an output ripple of about  $100mV_{P-P}$  at 71V with a 1mA load current. That level can be improved to less than 20mV by placing a low-cost electrolytic capacitor ( $10\mu F$ , 100V Nichicon VX-series) in parallel with a  $1\mu F$  ceramic capacitor (**Figure 4**). Additional filtering may be required to reduce the noise to lower levels. The typical noise level for an APD power supply is approximately 2mV. That level is easily achieved with a simple LC filter, given the MAX5026's fixed 500kHz switching frequency.

**Figure 5**'s schematic is an APD power supply with digitally adjustable output voltage. A  $\mu C$  in the control loop reads the thermistor value, provides temperature compensation, corrects the thermistor curvature through lookup tables, and compensates for gain variations due to APD manufacturing. In this application circuit, the 10-bit DAC (U2) provides approximately 45mV resolution when varying the output voltage from 25V to 71V.




Figure 4. This graph demonstrates the output voltage ripple for Figure 1's circuit with  $V_{OUT} = 71V$ ,  $I_{OUT} = 1mA$ , and a  $10\mu F$  electrolytic capacitor in parallel with the  $1\mu F$  output capacitor. The vertical axis is 50mV/div and the horizontal axis is  $200\mu s/div$ .

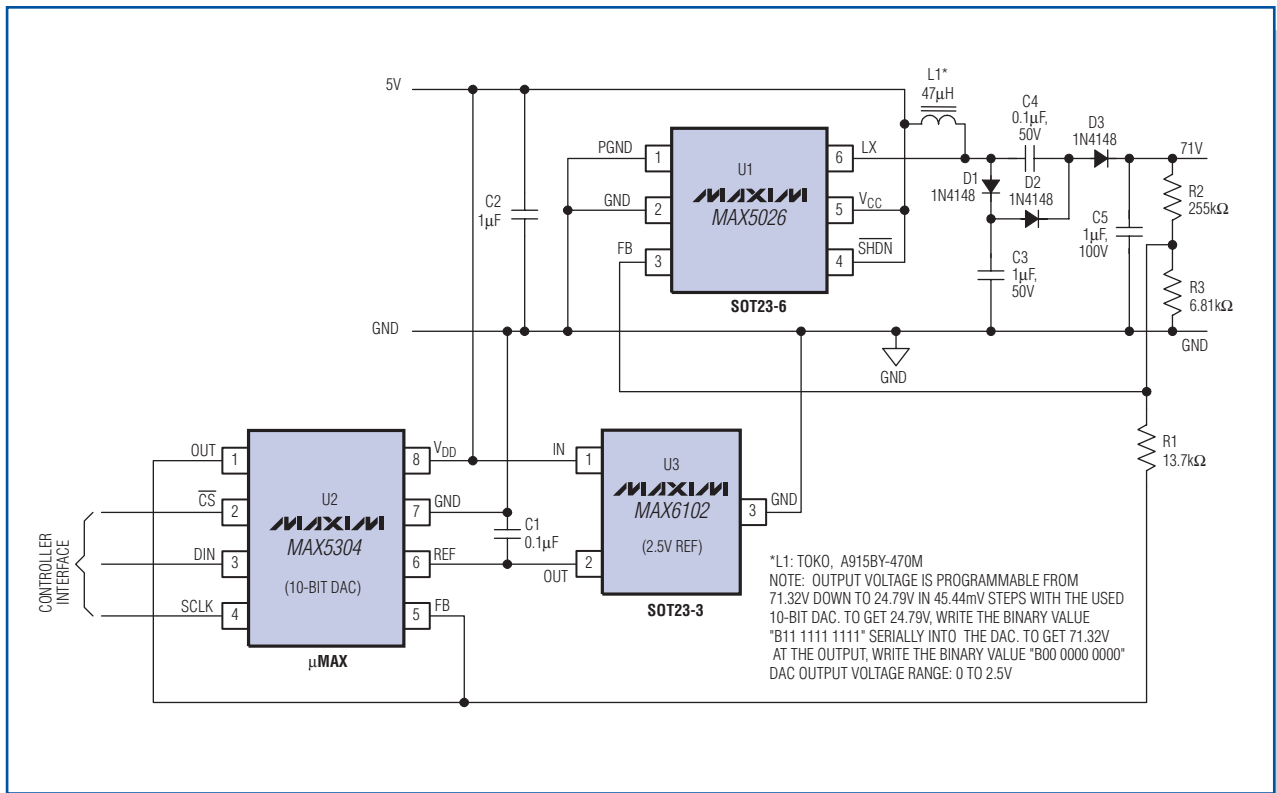


Figure 5. The output voltage in this low-noise APD bias power supply is digitally programmable from 25V to 71V in 45mV increments.

# Understanding successive approximation ADCs

Successive approximation register (SAR) analog-to-digital converters (ADCs) are frequently the architecture of choice for medium-to-high-resolution applications with sample rates under 5MSPs. SAR ADCs most commonly range in resolution from 8 bits to 16 bits and provide low power consumption as well as a small form factor. This combination makes them ideal for a wide variety of applications, such as portable/battery-powered instruments, pen digitizers, industrial controls, and data/signal acquisition.

As the name implies, the SAR ADC basically implements a binary search algorithm. Therefore, while the internal circuitry may be running at several megahertz (MHz), the ADC sample rate is a fraction of that number due to the successive approximation algorithm.

## SAR ADC architecture

Although there are many variations in implementing a SAR ADC, the basic architecture is quite simple (Figure 1). The analog input voltage ( $V_{IN}$ ) is held on a track/hold. To implement the binary search algorithm, the N-bit register is first set to midscale (that is, 100...00, where the MSB is set to '1'). This forces the digital-to-analog converter (DAC) output ( $V_{DAC}$ ) to be  $V_{REF}/2$ , where  $V_{REF}$  is the reference voltage provided to the ADC. A comparison is then performed to determine if

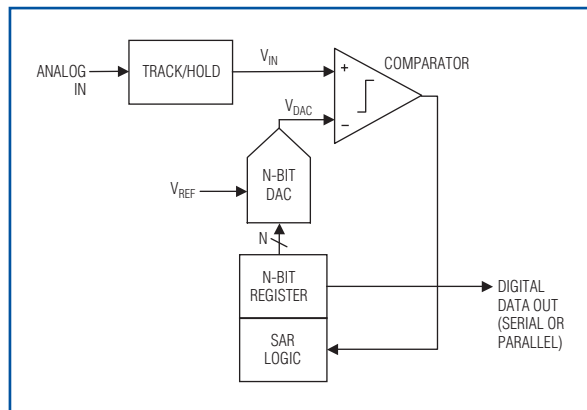


Figure 1. This diagram shows a simplified N-bit SAR ADC architecture.

$V_{IN}$  is less than or greater than  $V_{DAC}$ . If  $V_{IN} > V_{DAC}$ , the comparator output is a logic high or '1' and the MSB of the N-bit register remains at '1'. Conversely, if  $V_{IN} < V_{DAC}$ , the comparator output is a logic low and the MSB of the register is cleared to logic '0'. The SAR control logic then moves to the next bit down, forces that bit high, and does another comparison. The sequence continues all the way down to the LSB. Once this is done, the conversion is complete, and the N-bit digital word is available in the register.

Figure 2 shows an example of a 4-bit conversion. The y-axis (and the bold line in the figure) represents the DAC output voltage. In the example, the first comparison shows that  $V_{IN} < V_{DAC}$ . Thus, bit 3 is set to '0'. The DAC is then set to 0100<sub>2</sub> and the second comparison is performed. As  $V_{IN} > V_{DAC}$ , bit 2 remains at '1'. The DAC is then set to 0110<sub>2</sub>, and the third comparison is performed. Bit 1 is set to '0' and the DAC is then set to 0101<sub>2</sub> for the final comparison. Finally, bit 0 remains at '1' because  $V_{IN} > V_{DAC}$ .




Figure 2. This graph shows the DAC output voltage for successive decisions in a 4-bit SAR architecture.

Notice that four comparison periods are required for a 4-bit ADC. Generally speaking, an N-bit SAR ADC will require N comparison periods and will not be ready for the next conversion until the current one is complete. This explains why these ADCs are power- and space-efficient, yet are rarely seen in speed-and-resolution combinations beyond a few MSPs at 14 bits to 16 bits. Some of the smallest ADCs available on the market are based on the SAR architecture. The QSPI™ serially interfaced MAX1115–MAX1118 series of 8-bit ADCs as well as

QSPI is a trademark of Motorola, Inc.

their higher resolution counterparts (the 10-bit MAX1086 and the 12-bit MAX1286) fit in tiny SOT23 packages measuring 3mm x 3mm. The I<sup>2</sup>C™-compatible MAX1036/MAX1037 squeeze four 8-bit ADC channels and a reference into a SOT23 package.

One other feature of SAR ADCs is that the power dissipation scales with the sample rate unlike flash or pipelined ADCs, which usually have constant power dissipation versus sample rate. This is especially useful in low-power applications or applications where the data acquisition is not continuous (for example, PDA digitizers like the MAX1233).

### In-depth SAR analysis

Two critical components of a SAR ADC are the comparator and the DAC. As we shall see later, the track/hold shown in Figure 1 can be embedded in the DAC and may not be an explicit circuit.

A SAR ADC's speed is limited by:

- The DAC's settling time, which must settle to within the resolution of the overall converter (for example, 1/2LSB).
- The comparator, which must resolve small differences in  $V_{IN}$  and  $V_{DAC}$  within the specified time.
- The logic overhead.

### DACs

The DAC's maximum settling time is usually determined by the MSB settling time. This is simply because the MSB transition represents the largest excursion of the

DAC output. In addition, the linearity of the overall ADC is limited by the linearity of the DAC. Therefore, SAR ADCs with more than 12 bits of resolution often require some form of trimming or calibration to achieve the necessary linearity. This is due to the inherent component-matching limitations. Although it is somewhat process and design dependent, component matching limits the linearity to about 12 bits in practical DAC designs. Many SAR ADCs use a capacitive DAC that provides an inherent track/hold function. Capacitive DACs employ the principle of charge redistribution to generate an analog output voltage. Since these types of DACs are prevalent in SAR ADCs, it is beneficial to discuss their operation.

A capacitive DAC consists of an array of N capacitors with binary-weighted values plus one "dummy LSB" capacitor. **Figure 3** shows an example of a 16-bit capacitive DAC connected to a comparator. During the acquisition phase, the array's common terminal (the terminal at which all the capacitors share a connection) is connected to ground and all free terminals are connected to the input signal (Analog In or  $V_{IN}$ ).

After acquisition, the common terminal is disconnected from ground and the free terminals are disconnected from  $V_{IN}$ , effectively trapping a charge proportional to the input voltage on the capacitor array. The free terminals of all the capacitors are then connected to ground, driving the common terminal negative to a voltage equal to  $-V_{IN}$ .

As the first step in the binary search algorithm, the free terminal of the MSB capacitor is disconnected from ground and connected to  $V_{REF}$ , driving the common

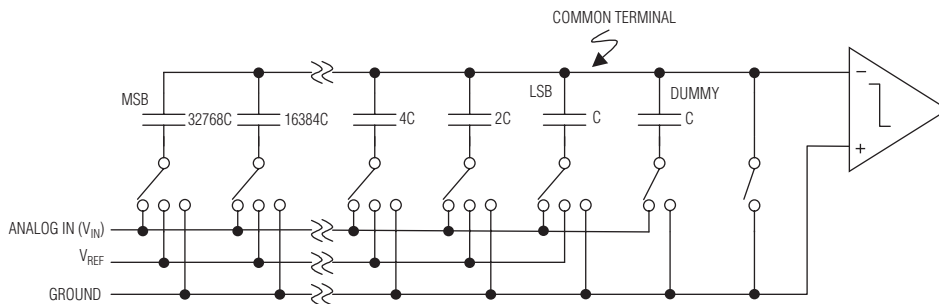


Figure 3. This diagram shows the basic architecture of a 16-bit capacitive DAC and illustrates the switch arrangement and the interface to the comparator.

I<sup>2</sup>C is a trademark of Phillips Corp.

terminal in the positive direction by an amount equal to  $1/2V_{REF}$ . For example, if  $V_{IN}$  is equal to  $3/4V_{REF}$ , connecting the MSB capacitor to  $V_{REF}$  and the rest of the capacitors to ground will drive the common terminal to  $(-3/4V_{REF} + 1/2V_{REF}) = -1/4V_{REF}$ . When this voltage is compared to ground, the comparator output yields a logic '1', implying that the MSB is greater than  $1/2V_{REF}$ .

Conversely, if  $V_{IN}$  is equal to  $1/4V_{REF}$ , the common terminal voltage is  $(-1/4V_{REF} + 1/2V_{REF}) = +1/4V_{REF}$ , and the comparator output is a logic '0'. Following this, the next-largest capacitor is disconnected from ground and connected to  $V_{REF}$ , and the comparator determines the next bit. This continues until all bits have been determined.

## DAC calibration

In an ideal DAC, each of the capacitors associated with the data bits should be exactly twice the value of the next smaller capacitor. In high-resolution ADCs (for example, a 16-bit ADC), this can result in a range of values too wide to be realized for an economically feasible size. A 16-bit SAR ADC such as the MAX195 utilizes a capacitor array that actually consists of two arrays that are capacitively coupled to reduce the LSB array's effective value. The capacitors in the MSB array are production-trimmed to reduce errors. Small variations in the LSB capacitors contribute insignificant errors to the 16-bit result. Unfortunately, trimming alone does not yield 16-bit performance or compensate for changes in performance due to changes in temperature, supply voltage, and other parameters.

For this reason, the MAX195 includes a calibration DAC for each capacitor in the MSB array. These DACs are capacitively coupled to the main DAC output and offset the main DAC's output according to the value on their digital inputs.

During calibration, the correct digital code to compensate for the error in each MSB capacitor is determined and stored. Thereafter, the stored code is provided to the appropriately calibrated DAC whenever the corresponding bit in the main DAC is high, compensating for errors in the associated capacitor. Calibration is usually initiated by the user or done automatically on power-up.

To reduce the effects of noise, each calibration experiment is performed many times (about 14,000 clock cycles in the MAX195) and the results are averaged. Calibration is best performed when the power-supply voltages are stable. High-resolution ADCs should be recalibrated any

time there is a significant change in supply voltage(s), temperature, reference voltage, or clock characteristics because these parameters affect the DC offset. If linearity is the only concern, much larger changes in these parameters can be tolerated. Since the calibration data is stored digitally, there is no need to perform frequent conversions to maintain accuracy.

## Comparators

Comparators require speed and accuracy. Although comparator offset does not affect overall linearity, it appears as an offset in the overall transfer characteristic. In addition, offset-cancellation techniques are usually applied to reduce the comparator offset. Noise, however, is a concern and the comparator is usually designed to have input-referred noise less than 1LSB. Additionally, the comparator needs to resolve voltages within the accuracy of the overall system; in other words, the comparator needs to be as accurate as the overall system.

## SAR ADCs versus other ADC architectures

### Pipelined ADCs

A pipelined ADC such as the MAX1200 employs a parallel structure in which each stage works concurrently on one to a few bits of successive samples. The inherent parallelism increases throughput, but at the expense of power consumption and latency.

Latency, in this case, is defined as the difference between the time an analog sample is acquired by the ADC and the time when the digital data is available at the output.

For instance, a five-stage pipelined ADC will have at least five clock cycles of latency, whereas a SAR has only one clock cycle of latency. Note that the latency definition applies only to the throughput of the ADC, not the internal clock of a SAR, which runs at many times the frequency of the throughput.

Pipelined ADCs frequently have digital error-correction logic to reduce the accuracy requirement of the flash ADCs (that is, comparators) in each pipeline stage. On the other hand, a SAR ADC requires the comparator to be as accurate as the overall system.

A pipelined ADC generally takes up significantly more silicon area than an equivalent SAR. Like a SAR, a pipelined ADC with more than 12 bits of accuracy usually requires some form of trimming or calibration.



## Flash ADCs

Flash ADCs such as the MAX117/MAX104 are comprised of large banks of comparators, each consisting of wideband, low-gain preamplifier(s) followed by a latch. The preamplifiers only have to provide gain but do not need to be linear or accurate; meaning, only the comparators' trip points have to be accurate. As a result, a flash ADC is the fastest architecture available. The primary trade-off for speed is the significantly lower power consumption and the smaller form factor. Extremely fast, 8-bit flash ADCs such as the MAX104/MAX106/MAX108 (and their folding/interpolation variants) have sampling rates as high as 1.5Gsps. It is much harder to find 10-bit flash ADCs, while 12-bit or higher flash ADCs are not commercially viable products. In a flash ADC, the number of comparators goes up by a factor of 2 for every extra bit of resolution, while at the same time, each comparator has to be twice as accurate. In a SAR ADC, however, the increased resolution requires more accurate components, yet the complexity does not increase exponentially. Of course, SAR ADCs are not capable of speeds at all comparable to those of flash ADCs.

## Sigma-delta converter ADCs

Traditional oversampling/sigma-delta converters commonly used in digital audio applications have limited bandwidths of about 22kHz. Recently, some high-bandwidth sigma-delta converters have reached bandwidths of 1MHz to 2MHz with 12 bits to 16 bits of resolution. These are usually very high-order, sigma-delta modulators (i.e. fourth order or higher) incorporating a multibit ADC and a multibit feedback DAC.

Sigma-delta converters such as the MAX1400/ MAX1403 have the innate advantage of requiring no special trimming or calibration, even to attain 16 bits to 18 bits of resolution. They also do not require anti-alias filters with

steep roll-offs at the analog inputs, because the sampling rate is much higher than the effective bandwidth.

The oversampling nature of the sigma-delta converter may also tend to "average out" any system noise at the analog inputs. However, sigma-delta converters trade speed for resolution. The need to sample many times (at least 16 times and often more) to produce one final sample dictates that the internal analog components in the sigma-delta modulator operate much faster than the final data rate. Designing the digital decimation filter is also a challenge and consumes a lot of silicon area. In the near future, the fastest high-resolution sigma-delta converters are not expected to have significantly higher bandwidth than a few megahertz.

## Conclusion

In summary, the primary advantages of SAR ADCs are low power consumption, high resolution, high accuracy, no latency in output data, and a small form factor. Because of these benefits, SAR ADCs can often be integrated with other larger functions. The main limitations of SAR architecture are that the lower sampling rates and requirements for building blocks, such as the DAC and comparator, need to be as accurate as the overall system.

## References:

1. MAX195 data sheet; Rev. 1, 12/97, Maxim Integrated Products.
2. Razavi, Behzad; *Principles of Data Conversion System Design*; IEEE Press, 1995.
3. Van De Plassche, Rudy; *Integrated Analog-to-Digital and Digital-to-Analog Converters*; Kluwer Academic Publishers, 1994.
4. "Pipelined ADCs—Recent Advances," Application Note, 2001, Maxim Integrated Products.

# Calculating MOSFET power dissipation for high-power, portable DC-DC converters

Perhaps the toughest challenge that designers of portable-power supplies face is powering modern high-performance CPUs. Recently, their supply currents have doubled every two years. In fact, today's portable core supplies can require up to 40A or more, at between 0.9V and 1.75V. But while current requirements have increased steadily, the space available for power supplies has not—a fact that has stretched thermal designs to the limit and beyond.

Supplies with such high current are typically broken into two or more phases, with each phase handling between

15A and 25A. This approach eases component selection. For example, a 40A supply essentially becomes two 20A supplies. But because this approach does not create additional board space, it hardly eases the thermal design challenge.

MOSFETs are the most difficult components to specify for high-current power supplies. This is especially true for notebook computers, an environment where heatsinks, fans, heatpipes, and other means of disposing heat are typically reserved for the CPU itself. Thus, the power supply often contends with cramped space, still air, and heat from nearby components. Moreover, nothing is available to aid power dissipation except a minimal amount of PC board copper underneath the supply.

MOSFET selection begins by choosing devices that can handle the required current, given an adequate thermal dissipation path. It ends with quantifying the needed thermal dissipation and ensuring the dissipation path. This article provides step-by-step instructions for calculating the power dissipation of these MOSFETs and determining the temperature at which they operate. It then illustrates these concepts by stepping through the design of one 20A phase of a multiphase, synchronous-rectified, step-down CPU core supply.

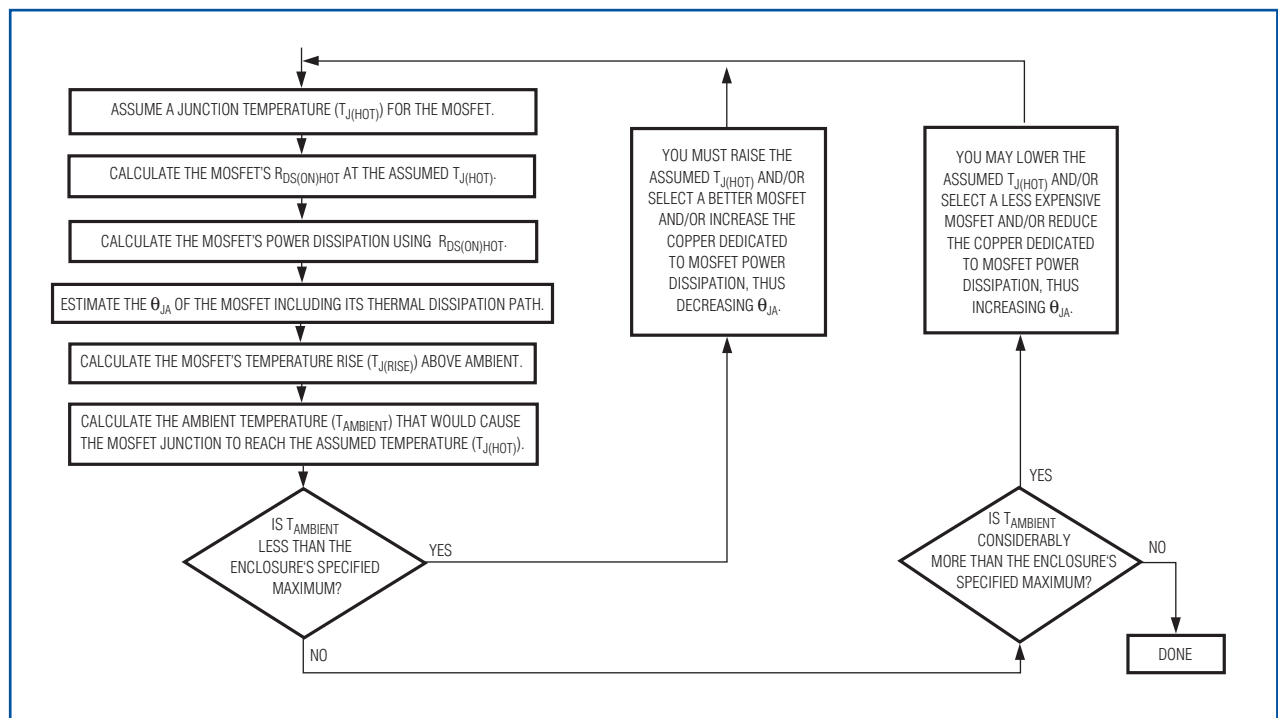


Figure 1. This flow chart represents the iterative process by which each of the MOSFETs (the synchronous rectifier and the switching MOSFET) is chosen. During this process, the junction temperature of each MOSFET is assumed, and both the MOSFET's power dissipation and allowable ambient temperature are calculated. The process ends when the allowable ambient temperature is at, or slightly above the maximum temperature expected within the enclosure that houses the power supply and the circuitry it powers.

## Calculating MOSFET power dissipation

To determine whether or not a MOSFET is suitable for a particular application, you must calculate its power dissipation, which consists mainly of resistive and switching losses:

$$PD_{\text{device total}} = PD_{\text{resistive}} + PD_{\text{switching}}$$

Because a MOSFET's power dissipation depends greatly on its on-resistance ( $R_{\text{DS(ON)}}$ ), calculating  $R_{\text{DS(ON)}}$  seems a good place to start. But a MOSFET's  $R_{\text{DS(ON)}}$  depends on its junction temperature ( $T_{\text{J}}$ ). In turn,  $T_{\text{J}}$  depends on both the power dissipated in the MOSFET and the thermal resistance ( $\theta_{\text{JA}}$ ) of the MOSFET. So, it is hard to know where to begin. Since several terms within the power dissipation calculation depend on each other, an iterative process is useful to determine this number (**Figure 1**).

The iterative process starts by first assuming a junction temperature for each MOSFET, then calculating each MOSFET's individual power dissipation and allowable ambient temperature. The process ends when the allowable ambient air temperature is at or slightly above the expected maximum temperature within the enclosure that houses the power supply and other circuitry it powers.

Making this calculated ambient temperature as high as possible may be tempting, but it is not usually a good idea. Doing so requires a more expensive MOSFET, more copper underneath the MOSFET, or moving more air by a larger, faster fan—all of which is unwarranted.

In a sense, assuming a MOSFET junction temperature and then calculating an associated ambient temperature entails working backwards. After all, the ambient temperature determines the MOSFET's junction temperature—not the other way around. But the calculations required when starting with an assumed junction temperature are easier to accomplish than when starting with an assumed ambient temperature and working from there.

For both the switching MOSFET and the synchronous rectifier, select a maximum permitted die junction temperature ( $T_{\text{J(HOT)}}$ ) to use as a starting point for this iterative process. Most MOSFET data sheets only specify a maximum  $R_{\text{DS(ON)}}$  at +25°C. But recently, some have offered maximums at +125°C as well. MOSFET  $R_{\text{DS(ON)}}$  increases with temperature, exhibiting typical temperature coefficients that range from 0.35%/°C to 0.5%/°C (**Figure 2**). If in doubt, use the more pessimistic tempera-

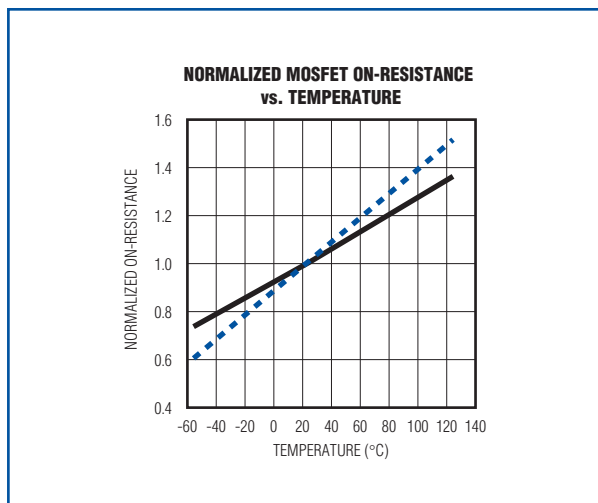


Figure 2. Typical power MOSFET on-resistance temperature coefficients range from 0.35% per degree (solid line) to 0.5% per degree (dashed line).

ture coefficient and the MOSFET's +25°C specification (or its +125°C specification, if available) to calculate an approximate maximum  $R_{\text{DS(ON)}}$  at your chosen  $T_{\text{J(HOT)}}$ :

$$R_{\text{DS(ON)HOT}} = R_{\text{DS(ON)SPEC}} \times [1 + 0.005 \times (T_{\text{J(HOT)}} - T_{\text{SPEC}})]$$

where  $R_{\text{DS(ON)SPEC}}$  is the MOSFET on-resistance used for the calculation, while  $T_{\text{SPEC}}$  is the temperature at which  $R_{\text{DS(ON)SPEC}}$  is specified. Use the calculated  $R_{\text{DS(ON)HOT}}$  to determine the power dissipation of both the synchronous rectifier and switching MOSFETs as described below.

The following paragraphs discuss calculating each MOSFET's power dissipation at its assumed die temperature, followed by the additional steps to complete this iterative process (the entire procedure is detailed in Figure 1).

### The synchronous rectifier's power dissipation

For all but the lightest loads, the drain-to-source voltage of the synchronous rectifier's MOSFET is clamped by the catch diode during turn-on and turn-off. Therefore, the synchronous rectifier incurs no switching losses, making its power dissipation easy to calculate. Only resistive losses must be considered.

The worst-case losses occur at the maximum duty factor of the synchronous rectifier, which occurs when the input voltage is at its maximum. Through the use of the

synchronous rectifier's  $R_{DS(ON)HOT}$  and its duty factor, along with Ohm's Law, you can calculate its approximate power dissipation:

$$PD_{\text{synchronous rectifier}} = [I_{LOAD}^2 \times R_{DS(ON)HOT}] \times [1 - (\frac{V_{OUT}}{V_{IN(MAX)}})]$$

## The switching MOSFET's power dissipation

The switching MOSFET's resistive losses are calculated much like the synchronous rectifier's, using its (different) duty factor and  $R_{DS(ON)HOT}$ :

$$PD_{\text{resistive}} = [I_{LOAD}^2 \times R_{DS(ON)HOT}] \times (\frac{V_{OUT}}{V_{IN}})$$

Calculating the switching MOSFET's switching loss is difficult since it depends on many hard to quantify and typically unspecified factors that influence both turn-on and turn-off. Use the rough approximation in the following formula as the first step in evaluating a MOSFET and verify performance on the lab bench:

$$PD_{\text{switching}} = \frac{C_{RSS} \times V_{IN}^2 \times f_{sw} \times I_{LOAD}}{I_{GATE}}$$

where  $C_{RSS}$  is the MOSFET's reverse-transfer capacitance (a data sheet parameter),  $f_{sw}$  is the switching frequency, and  $I_{GATE}$  is the MOSFET gate-driver's sink/source current at the MOSFET's turn-on threshold (the  $V_{GS}$  of the gate-charge curve's flat portion).

Once you have narrowed the choice to a specific generation of MOSFETs based on cost (the cost of a MOSFET is very much a function of the specific generation to which it belongs), select the device within the generation that will minimize power dissipation. This is the one with equal resistive and switching losses. Using a smaller (faster) device increases resistive losses more than it decreases switching losses, and a larger (low  $R_{DS(ON)}$ ) device increases switching losses more than it decreases resistive losses.

If  $V_{IN}$  varies, calculate the switching MOSFET's power dissipation at both  $V_{IN(MAX)}$  and  $V_{IN(MIN)}$ . The MOSFET's worst-case power dissipation will occur at either the minimum or the maximum input-voltage level. The dissipation is the sum of two functions: the resistive dissipation, which is highest at  $V_{IN(MIN)}$  (the higher duty factor), and the switching dissipation, which is highest at  $V_{IN(MAX)}$  (because of the  $V_{IN}^2$  term). An optimal

selection has roughly equal dissipations at the  $V_{IN}$  extremes, balancing resistive and switching dissipations across the  $V_{IN}$  range.

If the dissipation at  $V_{IN(MIN)}$  is significantly higher, resistive losses dominate. In that case, consider a larger switching MOSFET (or more than one in parallel) to lower  $R_{DS(ON)}$ . But if the losses at  $V_{IN(MAX)}$  are significantly higher, consider decreasing the size of the switching MOSFET (or removing a MOSFET if multiple devices are used) to allow it to switch faster.

If the resistive and switching losses balance, but are still too high, there are several ways to proceed:

- Change the problem definition. For example, redefine the input voltage range.
- Change the switching frequency to lower switching losses, possibly allowing a larger and lower  $R_{DS(ON)}$  switching MOSFET.
- Increase the gate-driver current, possibly lowering switching losses. The MOSFET's own internal gate resistance, which ultimately limits the gate-driver current, places a practical limit on this approach.
- Use an improved MOSFET technology that might simultaneously switch faster, have lower  $R_{DS(ON)}$ , and have lower gate resistance.

Fine-tuning the MOSFET's size beyond a certain point may not be possible due to a limited choice of devices. The bottom line is that the MOSFET's worst-case power must be dissipated.

## Thermal resistance

Refer to Figure 1 for the next step in the iterative process to determine the right MOSFETs for both the synchronous rectifier and the switching MOSFET. This step is the calculation of the ambient air temperature surrounding each MOSFET that would cause the assumed MOSFET junction temperature to be reached. To do this, first determine the junction-to-ambient thermal resistance ( $\theta_{JA}$ ) of each MOSFET.

Thermal resistance can be difficult to estimate. While it is relatively easy to measure the  $\theta_{JA}$  of a single device on a simple PC board, it can be hard to predict thermal performance in an actual power supply within a system, where many heat sources compete for limited dissipation paths. If multiple MOSFETs are used in parallel, you can calculate their combined thermal resistance in the same way as the equivalent resistance of two or more paralleled resistors.

Start with the MOSFET's  $\theta_{JA}$  specification. For single-die, 8-pin MOSFET packages,  $\theta_{JA}$  is usually near 62°C/W. For other packages, with thermal tabs or exposed heat slugs, it may range between 40°C/W to 50°C/W (**Table 1**). To calculate the MOSFET's die temperature rise above ambient, use the following equation:

$$T_{J(RISE)} = PD_{device\ total} \times \theta_{JA}$$

Next, calculate the ambient temperature that will cause the die to reach the assumed  $T_{J(HOT)}$ :

$$T_{AMBIENT} = T_{J(HOT)} - T_{J(RISE)}$$

If the calculated  $T_{AMBIENT}$  is lower than the enclosure's maximum specified ambient temperature (meaning the enclosure's maximum specified ambient temperature will cause the MOSFET's assumed  $T_{J(HOT)}$  to be exceeded), you must do one or more of the following:

- Raise the assumed  $T_{J(HOT)}$ , but not above the data sheet maximum.
- Lower the MOSFET power dissipation by choosing a more suitable MOSFET.
- Decrease  $\theta_{JA}$  by increasing the airflow or the amount of copper around the MOSFET.

Recalculate  $T_{AMBIENT}$  (employing a spreadsheet simplifies the multiple iterations typically required to select an acceptable design).

On the other hand, if the calculated  $T_{AMBIENT}$  is higher than the enclosure's maximum specified ambient temperature by a fair amount, any or all of the following optional steps can be taken:

- Lower the assumed  $T_{J(HOT)}$ .
- Reduce the copper dedicated to the MOSFET's power dissipation.
- Use a less expensive MOSFET.

These steps are optional since in this case, the MOSFET will not be damaged by excessive temperature. However, these steps reduce both board area and cost as long as the calculated  $T_{AMBIENT}$  remains higher than the enclosure's maximum temperature by some margin.

The biggest source of inaccuracy in this procedure is  $\theta_{JA}$ . Carefully read any data sheet notes associated with a  $\theta_{JA}$  specification. Typical specifications assume a device mounted on 1in<sup>2</sup> of 2oz copper. The copper performs much of the power dissipation, and different amounts of copper change  $\theta_{JA}$  dramatically. For example, the  $\theta_{JA}$  of a D-Pak might be 50°C/W with 1in<sup>2</sup> of copper. But with copper underlying just the package footprint, the  $\theta_{JA}$  more than doubles (Table 1).

With multiple MOSFETs in parallel, the  $\theta_{JA}$  mostly depends on the copper area to which they are mounted. The equivalent  $\theta_{JA}$  for two devices can be half that of one device, but only if the copper area is also doubled. That is, adding a parallel MOSFET without additional copper halves the  $R_{DS(ON)}$  but changes  $\theta_{JA}$  much less.

**Table 1. Typical thermal resistances of MOSFET packages**

| PACKAGE  | $\theta_{JA}$ (°C/W)<br>MINIMUM<br>FOOTPRINT | $\theta_{JA}$ (°C/W)<br>1in <sup>2</sup> of 2oz<br>COPPER | $\theta_{JA}$ (°C/W) |
|--|--|---|----------------------|
| SOT23 (Thermally Enhanced)                     | 270  | 200   | 75                   |
| SOT89  | 160  | 70  | 35                   |
| SOT223   | 110  | 45  | 15                   |
| 8-Pin $\mu$ MAX/Micro8<br>(Thermally Enhanced) | 160  | 70  | 35                   |
| 8-Pin TSSOP                                    | 200  | 100   | 45                   |
| 8-Pin SO<br>(Thermally Enhanced)               | 125  | 62.5  | 25                   |
| D-PAK  | 110  | 50  | 3                    |
| D2-PAK   | 70   | 40  | 2                    |

*Note: Thermal resistances vary among individual devices in the same package type and among similar packages from different manufacturers depending on package mechanical characteristics, die size, and mounting and bonding method. Carefully consider thermal information in the MOSFET data sheet.*

Lastly,  $\theta_{JA}$  specifications assume that no other devices contribute heat to the copper dissipation area. At high currents, every component in the power path, even PC board copper, generates heat. To avoid overheating the MOSFETs, carefully estimate the  $\theta_{JA}$  that the physical situation can realistically achieve and consider the following:

- Study the selected MOSFET's available thermal information.
- Investigate whether or not there is space available for additional copper, heatsinks, and other devices.
- Determine if increasing air flow is feasible.
- See if other devices contribute significant heat to the assumed dissipation path.
- Estimate excess heating or cooling from nearby components and spaces.

### Design example

The CPU core supply shown in **Figure 3** delivers 1.3V at 40A. Two identical 20A power stages operating at 300kHz supply the 40A output current. The MAX1718 master controller drives one stage, while the MAX1897 slave controller drives the other. The supply's input range spans 8V to 20V, with the specified maximum ambient temperature of the enclosure at +60°C.

The synchronous rectifier comprises two IRF7822 MOSFETs in parallel, with a combined maximum  $R_{DS(ON)}$  of 3.25m $\Omega$  at room temperature and approximately 4.7m $\Omega$  at +115°C (the assumed  $T_{J(HOT)}$ ). With a maximum duty factor of 94%, a 20A load current, and a 4.7m $\Omega$  maximum  $R_{DS(ON)}$ , these paralleled MOSFETs dissipate about 1.8W. Supplied with 2in<sup>2</sup> of copper to dissipate that power, the overall  $\theta_{JA}$  should be about 31°C/W. The temperature rise of the combined MOSFETs will be approximately +55°C, so this design will work with an ambient temperature up to +60°C.

The switching MOSFET has two IRF7811W MOSFETs in parallel, with a combined maximum  $R_{DS(ON)}$  of 6m $\Omega$  at room temperature and approximately 8.7m $\Omega$  at +115°C (the assumed  $T_{J(HOT)}$ ). The combined  $C_{RSS}$  is 240pF. The MAX1718's and MAX1897's 1 $\Omega$  gate drivers deliver approximately 2A. At  $V_{IN} = 8V$ , the resistive losses are 0.57W, and the switching losses are approximately 0.05W. At 20V, the resistive losses are 0.23W and the switching losses are approximately 0.29W. Total losses at each operating point roughly balance and the worst-case total loss is 0.61W at the minimum  $V_{IN}$ .

Because this level of power dissipation is not high, we can provide under 0.5in<sup>2</sup> of copper area to these MOSFETs, achieving an overall  $\theta_{JA}$  of approximately 55°C/W. This enables operation up to an ambient temperature of +80°C with a +35°C temperature rise.

The copper areas in this example are required for the MOSFETs alone. If other devices dissipate heat into those areas, more copper area will likely be required. If space is not available for the additional copper, reduce the total power dissipation, spread the heat to areas of low dissipation, or use active means to remove heat.

### Conclusion

Thermal management is one of the most difficult areas of high-power portable design. This difficulty makes the iterative process outlined above necessary. Although this process should bring the board designer close to the final design, lab work must ultimately determine whether the design process was sufficiently accurate. Calculating the MOSFET's thermal properties and ensuring their dissipation paths, while checking those calculations in the lab, help guarantee a robust thermal design.

*Reprinted with permission by Electronic Design, Oct. 14, 2002. Copyright 2002, Penton Media, Inc.*

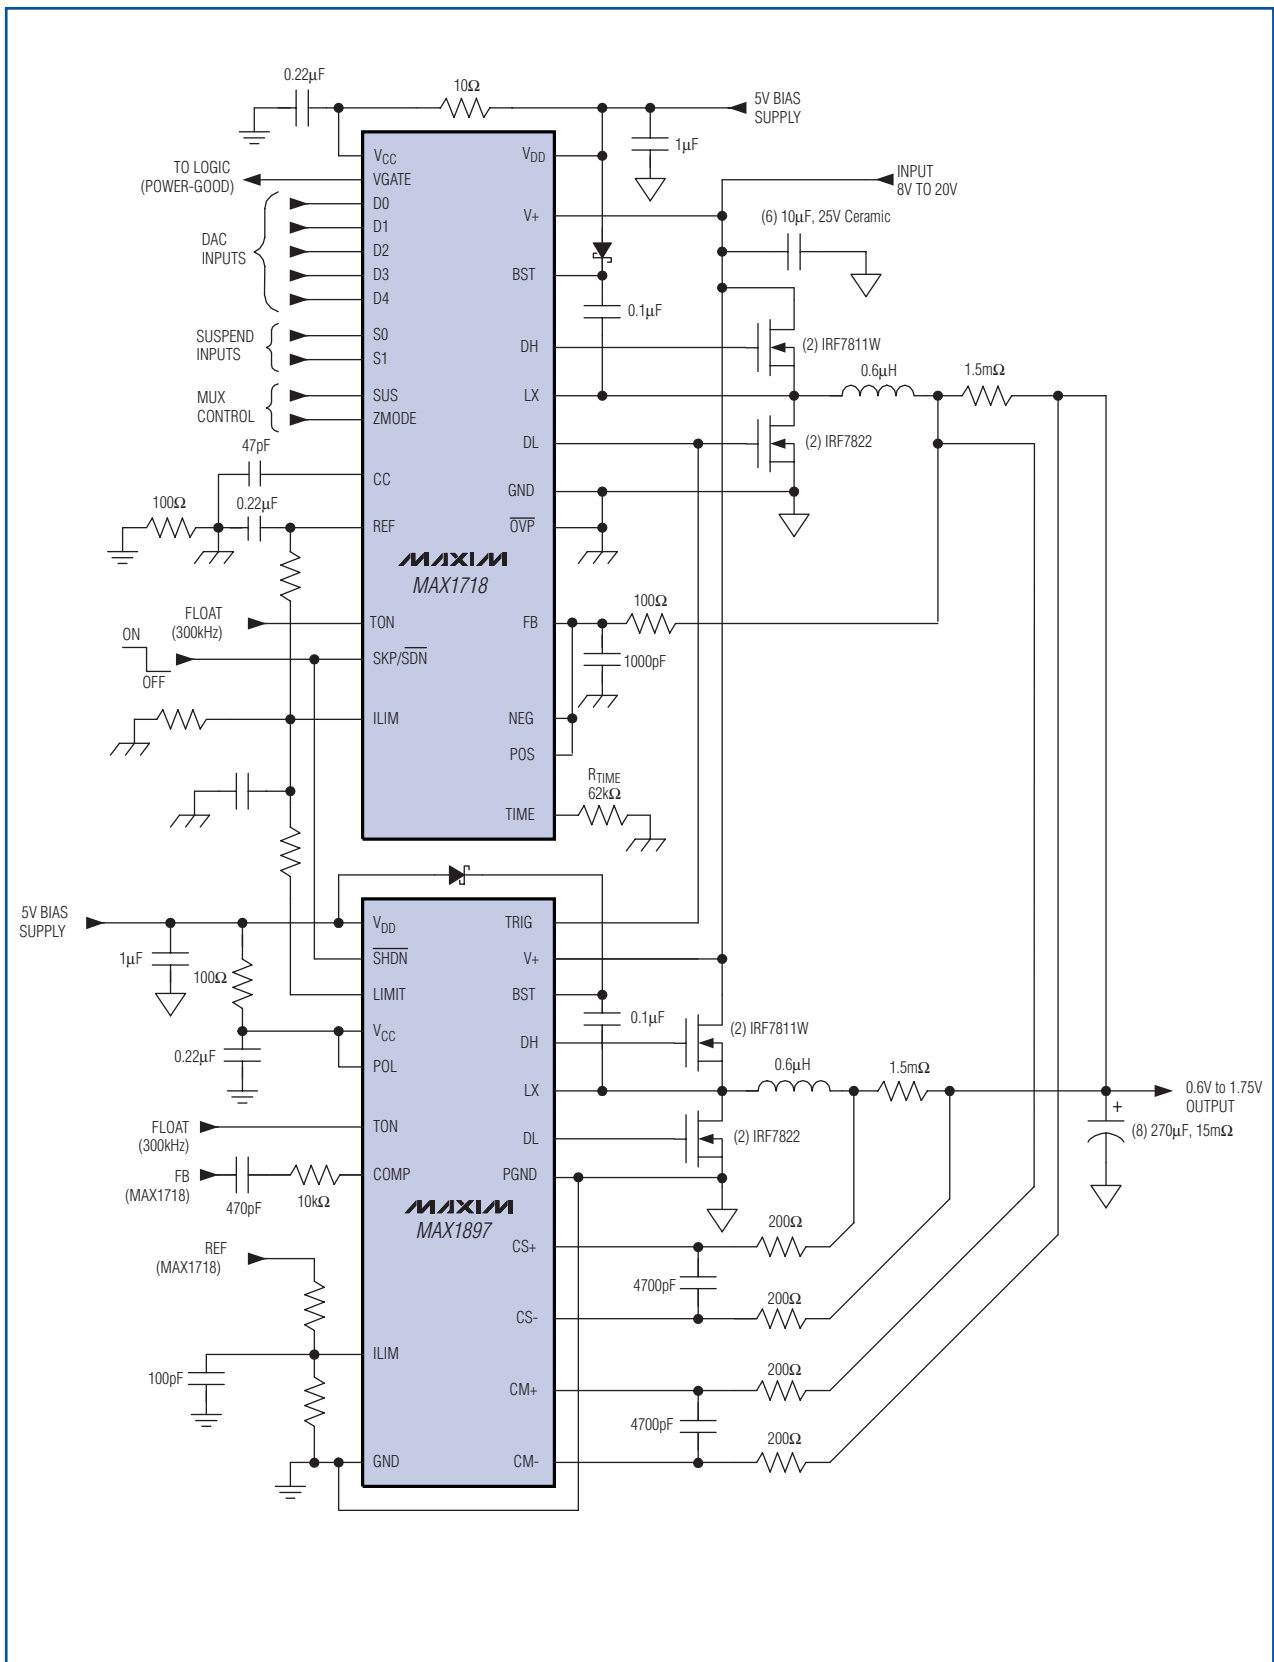


Figure 3. The MOSFETs for this step-down switching regulator were chosen using the iterative process described in this article. Board designers commonly use this type of switching regulator to power modern high-performance CPUs.

# DESIGN SHOWCASE

## Power meter achieves $\pm 1\%$ accuracy

Power meters provide an early warning of thermal overload by monitoring the power consumption in high-reliability systems. Power monitoring is especially suitable for systems in which the load voltage and current are both variable such as industrial heating and motor controllers. Such a power meter/controller (**Figure 1**) is based on the principle that power equals the product of current and voltage. Its typical accuracy is better than 1%.

A current sensor (U2) measures output current, and a four-quadrant analog voltage multiplier (U1 and U3) generates the product of output voltage and current. An optional unity-gain inverter (U4) inverts the inverted multiplier output. This power meter is most accurate for multiplier inputs (J1 and J2) between 3V and 15V. Choose the current-sense resistor,  $R_{SENSE}$ , as follows:

$$R_{SENSE}(\Omega) = \frac{1}{P(W)}$$

where  $R_{SENSE}$  is in ohms and P is the output power in watts. If power delivery to the load is 10W, for instance, choose  $R_{SENSE} = 0.1\Omega$ .

The Figure 1 circuit, with a  $0.1\Omega$  sense resistor, has a unity-gain transfer function in which the output voltage is proportional to load power. For instance, the output voltage is 10V when the load power is 10W. To change the gain transfer function, change the sense resistor as follows:

$$\text{Gain} = 10R_{SENSE}$$

**Figure 2** compares power-measurement error with load power for the Figure 1 circuit. Note that the accuracy is better than  $\pm 1\%$  for load power in the 3W to 14W range.

For proper operation, the analog multiplier must first be calibrated according to the following procedure (which also appears in Motorola's MC1495 data sheet). To

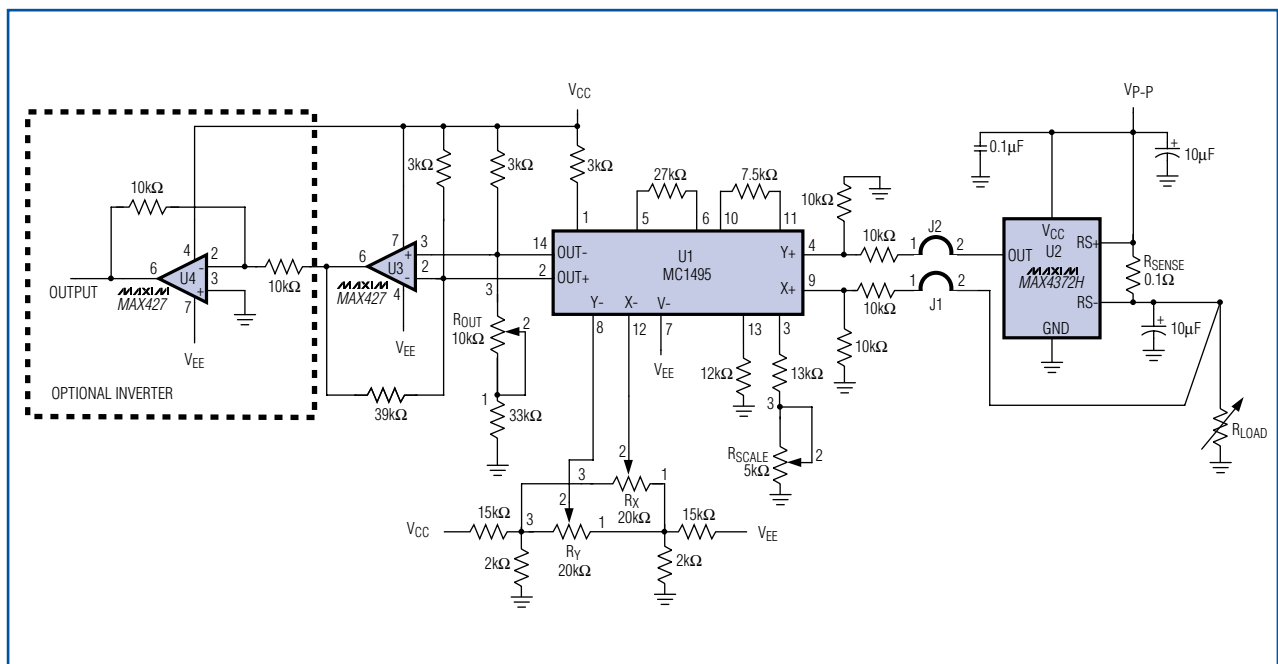


Figure 1. This power meter, whose output voltage is proportional to load power, achieves  $\pm 1\%$  accuracy.



calibrate the multiplier, remove jumper J1 (X input) and J2 (Y input), and use the following procedures:

- 1) **X-input offset adjustment:** Connect a 1.0kHz, 5V<sub>P-P</sub> sinewave to the Y input, and connect the X input to ground. Using an oscilloscope to monitor the output, adjust  $R_X$  for an AC null (zero amplitude) in the sinewave.
- 2) **Y-input offset adjustment:** Connect a 1.0kHz, 5V<sub>P-P</sub> sinewave to the X input, and connect the Y input to ground. Using an oscilloscope to monitor the output, adjust  $R_Y$  for AC null (zero amplitude) in the sinewave.
- 3) **Output offset adjustment:** Connect both X and Y inputs to ground. Adjust  $R_{OUT}$  until the output DC voltage is zero.
- 4) **Scale factor (Gain):** Connect both X and Y inputs to 10VDC. Adjust  $R_{SCALE}$  until the output voltage is 10VDC.
- 5) Repeat steps 1 through 4 as necessary.

*A similar idea appeared in EDN.*




Figure 2. This graph shows measured power has better than  $\pm 1\%$  accuracy for power levels between 3W and 14W.

# DESIGN SHOWCASE

## Portable devices derive 3.3V and 5V power supplies from USB ports

The universal serial bus (USB) port provides power in addition to its communication channel (D+, D-). While connected to the USB port for communications, battery-powered devices such as digital cameras, MP3 players, and PDAs can utilize USB power to charge their battery. **Figure 1**'s circuit utilizes USB power to

produce 3.3V and 5V supply rails and to charge a lithium (Li+) battery. U1 charges the battery, U2 steps up the battery voltage ( $V_{BATT}$ ) to 5V, and U3 steps down the 5V output to 3.3V.

U1, a Li+ battery charger, draws power from the USB port to charge the battery. Pulling its SETI terminal

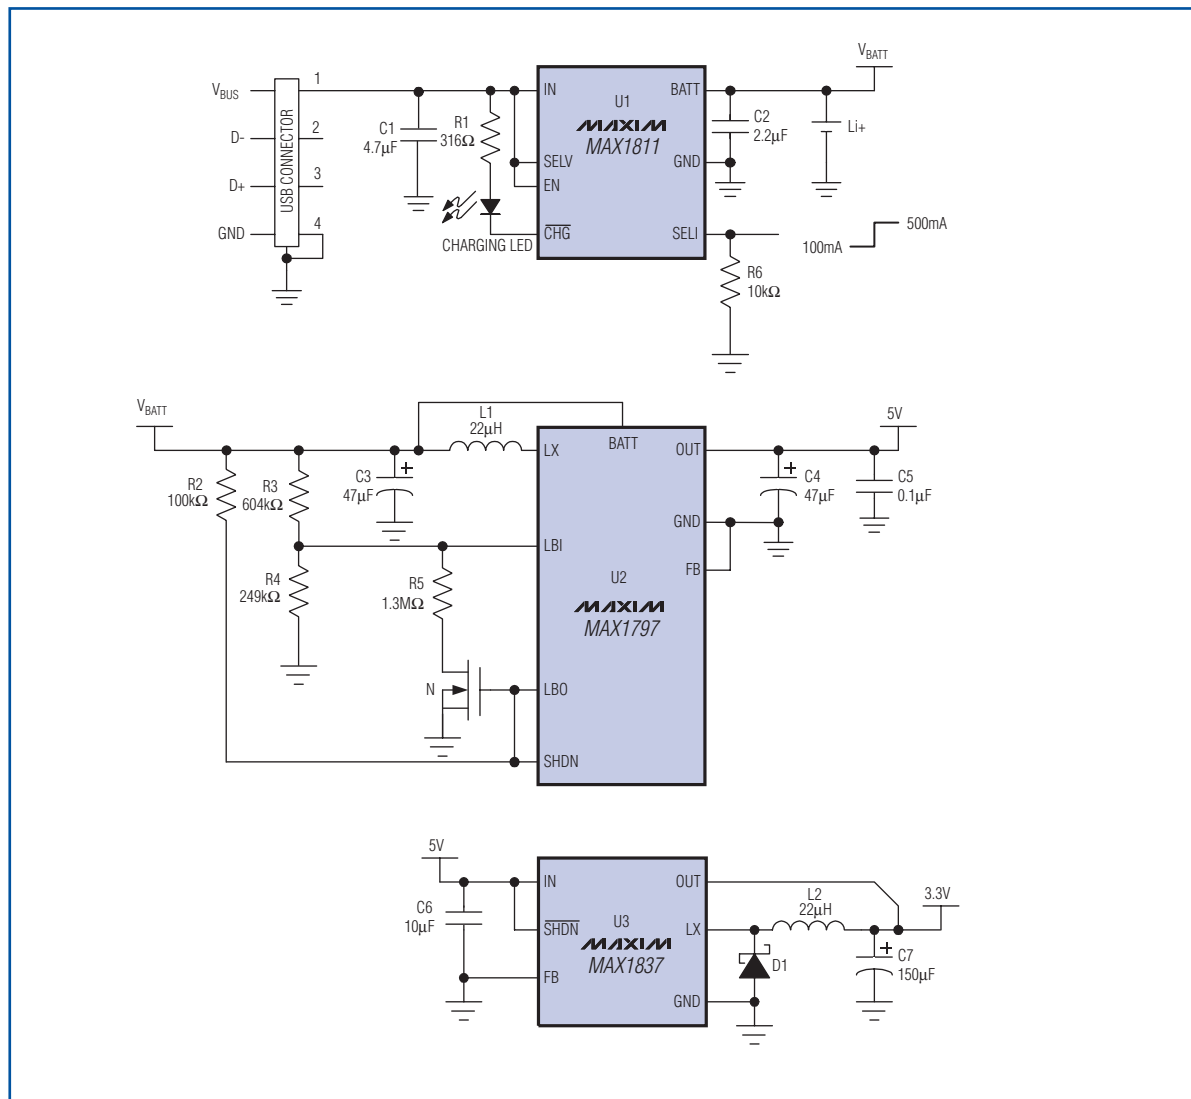


Figure 1. This circuit generates 5V and 3.3V supply voltages for portable applications while drawing power from the USB port.

low sets the charging current to 100mA for low-power USB ports, and pulling SETI high sets 500mA for high-power ports. Similarly, pulling SETV high or low configures the chip for charging a 4.2V or 4.1V Li+ battery. To protect the battery, U1's final charging voltage exhibits 0.5% accuracy. The  $\overline{\text{CHG}}$  terminal allows the chip to illuminate an LED during charging.

U2 is a step-up DC-DC converter that boosts  $V_{\text{BATT}}$  to 5V and delivers up to 450mA of current to its load. Its low-battery detection circuitry and true shutdown capability protect the Li+ battery. By disconnecting the battery from the output, "true" shutdown limits battery current to less than 2 $\mu$ A. The low-battery trip point is set by an external resistive divider between  $V_{\text{BATT}}$  and GND, connected to LBI. Connecting the low-battery output (LBO) to shutdown (SHDN) causes U2 to disconnect its load in response to a low-battery voltage.

The internal source impedance of a Li+ battery makes U2 susceptible to oscillation when its low-battery detection circuitry disconnects a low-voltage battery from its load. As the voltage drop across the battery's internal resistance is removed, the battery voltage increases and turns U2 back on. For example, a Li+ battery with 500m $\Omega$  internal resistance (while sourcing 500mA) drops 250mV across its internal resistance. When the U2 circuitry disconnects the load, the battery current drops to zero and the battery voltage increases by 250mV.

The N-channel FET at LBO eliminates this oscillation by adding hysteresis to the low-battery detection circuitry. Figure 1's circuit is configured for a low-battery trip voltage of 2.9V. When  $V_{\text{BATT}}$  drops

### USB POWER

- A low-power USB port provides 4.4V to 5.25V at 100mA.
- A high-power USB port provides 4.75V to 5.25V at 500mA.
- Due to voltage drops across USB cables and connectors, a USB device must be able to operate with 4.35V.
- A USB device must ensure that its maximum current consumption is 100mA, until configured for high power through software.

below 2.9V, LBO opens and allows SHDN to be pulled high, turning on the FET. With the FET turned on, the parallel combination of 1.3M $\Omega$  and 249k $\Omega$  eliminates oscillation by setting the battery turn-on voltage to 3.3V.

$$V_{\text{BATT(TURN-OFF)}} = V_{\text{LBI}} \times \frac{R3 + R4}{R4}$$

where  $V_{\text{LBI}} = 0.85\text{V}$ , and

$$V_{\text{BATT(TURN-ON)}} = V_{\text{LBI}} \times \frac{R3 + R'4}{R'4}$$

where

$$R'4 = \frac{R4 \times R5}{R4 + R5}$$

Finally, a step-down converter (U3) bucks 5V to 3.3V, and delivers up to 250mA to its load with efficiency exceeding 90%.

*A similar idea appeared in EDN.*