Managing Multiple Supply Voltages

By **Joe Chong**, Business Manager, Standard Products, and **Jay Scolio**, Technical Editorial Manager, Maxim Integrated Products, Sunnyvale, Calif.

Depending on system power requirements, power-supply monitoring and control may be accomplished using a variety of components ranging from power-on resets and voltage monitors to programmable system-management ICs.

For most electronic systems, monitoring system voltages with a power-on reset (POR) ensures proper initialization at power up. Furthermore, detecting brownout conditions with a POR minimizes possible code-execution problems that can corrupt memory or cause the system to execute improperly.

To improve reliability in high-end systems, it is often necessary to sequence a system's power supplies in the correct order to prevent its microcontrollers, microprocessors, DSPs or ASICs from latching up. In most cases, one or more microprocessor supervisor ICs can readily perform these sequencing and monitoring functions.

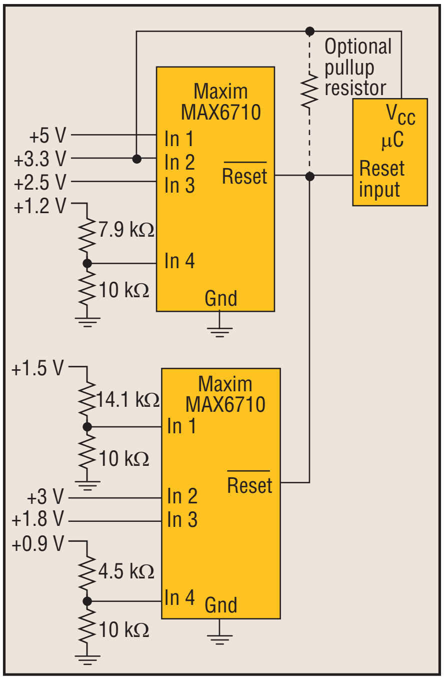

Fig. 1. *Two multivoltage supervisors with open-drain outputs monitor eight voltages and provide a single reset output.*

A related power management function, margining, is required in many applications to test a system's vulnerability to supply voltage fluctuations. As with sequencing, margining can be implemented using just a few added components. However, as supply voltages proliferate and power management requirements become more complex, fully integrated and programmable system-management devices offer a flexible alternative to managing supply monitoring, sequencing and margining.

Monitoring Voltages with Detectors and Power-On Resets

A simple way to monitor a system supply voltage is with a voltage detector, an IC that combines a comparator and an internal reference. When the supply voltage falls below the voltage detector's threshold, its output asserts to notify the system's microcontroller of an impending power failure. This warning gives the microcontroller an opportunity to back up its memory, turn on or off power supplies, and shut down the system in a controlled manner.

When a voltage detector changes states during power up or down, its output asserts after a short propagation delay.

When a voltage detector changes states during power up or down, its output asserts after a short propagation delay. This is fine for a power-fail warning. In most cases, a microcontroller's reset input requires a longer delay; during power-up, the system clocks and power supplies must stabilize and the processor's registers must initialize before releasing the microcontroller from its reset state.

A power-on reset (POR) is a form of microprocessor supervisor IC that provides this longer delay (called a reset timeout) to permit the system to fully initialize prior to allowing the microcontroller to operate. Also, when the supply voltage falls temporarily below the POR's threshold after power-up has occurred, this same delay applies after the supply voltage returns above the POR threshold. Power-on resets are available with several different fixed timeout periods and threshold voltages, some of which provide capacitor-adjustable timeout periods.

Most systems monitor the 3.3-V I/O logic supply. For systems requiring higher reliability, it may be necessary to monitor additional supplies, such as those that power cores and memory. Numerous multivoltage microprocessor supervisors are capable of performing this task, but the specific requirements of a given system can quickly reduce the number of choices.

Although most supervisors monitor standard voltages, such as 5 V, 3.3 V, 2.5 V, and 1.8 V, it is often necessary to monitor additional voltages because various components (e.g., memories, PLDs and ASICs) have unique power-supply requirements. As a result, you must decide whether

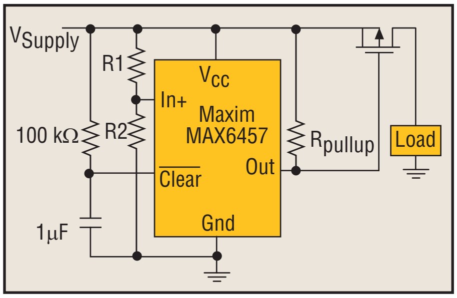

Fig. 2. *When this supervisory circuit detects an overvoltage condition, the P-channel MOSFET disconnects the supply.*

to use a fixed-threshold device (which requires no external resistors) or a more flexible adjustable-threshold device that accommodates changes as needed (but that requires external resistors). A device with a combination of fixed and adjustable thresholds might provide the best solution.

When selecting a device, it's important to choose one with a reference whose voltage is low enough to monitor the system's lowest voltage. When working with 0.8-V, 0.9-V and 1-V supplies, for example, a device with a standard 1.2-V reference won't work.

The number of supply voltages present in high-reliability systems has increased in recent years; 10 or more voltages are common. When monitoring a large number of voltages, you can end up using several supervisor devices. Multivoltage supervisors with open-drain outputs often are advantageous in these situations because their outputs can be OR'd together to provide a single output (Fig. 1).

It may be necessary to monitor some power supplies not just for undervoltage, but for overvoltage conditions as well. Overvoltage monitoring has become necessary in many systems to prevent damage to expensive processors and ASICs. A window detector, which monitors both overvoltage and undervoltage conditions, can be constructed with two voltage detectors and a reference. Alternatively, you can use a dedicated window detector IC. Another type of voltage protection circuit includes an external P-channel MOSFET that shuts off a supply if the supply voltage exceeds a specified level (Fig. 2).

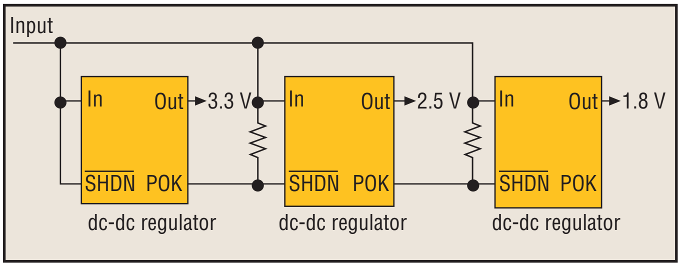

Fig. 3. *A power supply with a POK output provides a convenient way to sequence other supplies.*

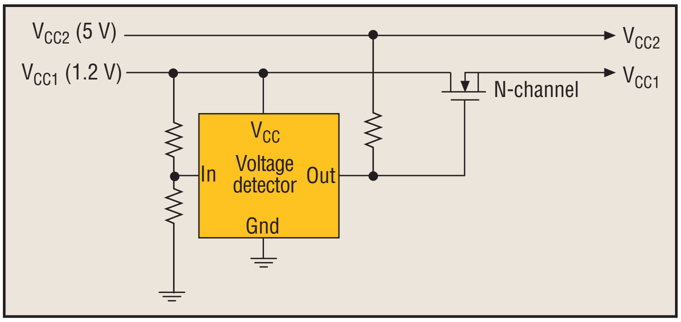

Fig. 4. *If a higher voltage is available, a voltage detector can sequence a low voltage by turning on an N-channel MOSFET.*

Sequencing Power Supplies

You can conveniently sequence power by using the enable or shutdown pins of the power regulators. Under this "daisy-chaining" scheme, when a power supply first comes up, it asserts its Power OK (POK) signal (if it has one) to notify other circuitry that its voltage is within tolerance. The POK output connects to the shutdown or enable pin of the second regulator and turns on that regulator when it goes active (Fig. 3). For those situations where a longer delay is needed, some regulators include a POR; this arrangement allows a longer time delay before turning on the next power supply in the sequence.

When a POK signal isn't available, you can monitor a power supply's output with a voltage detector or a POR by connecting the detector or POR output to the second supply's shutdown or enable input. The second supply turns on when the monitored voltage exceeds a specific threshold.

When used with noisy power supplies, a voltage detector might unnecessarily turn a regulator on and off several times, especially if the monitored voltage level is near its trip threshold. In these situations, a power-on reset circuit can minimize this effect—a benefit of the POR's timeout period.

When the monitored voltage falls below the supervisor's threshold, the POR's output asserts and remains asserted for at least the minimum reset timeout period after the monitored voltage returns above the threshold. The voltage must be above the reset threshold continuously during the timeout period in order for the supervisor to de-assert, thus preventing the power supply from cycling repeatedly.

In addition to the benefits listed above, using a POR to generate a signal for the shutdown or enable pin allows you to control the turn-on time; PORs have reset timeouts that range from a few microseconds to more than one second. Also, capacitor-adjustable devices allow you to change the timeout period of a given device.

A POR circuit also gives you the ability to control other power-up sequencing situations. For instance, in a system with three power supplies, you might want the first two supplies to be valid before the third supply is activated. If a single regulator without a POK output generates the first two supplies, you can use a dual-voltage POR to monitor its two voltages. This POR's output then controls the sequencing of the third supply by feeding its enable or shutdown pin.

To sequence larger numbers of supplies, you can use multivoltage devices. For example, a quad voltage detector is suitable for sequencing four voltages. Furthermore, devices with multiple reset outputs with different delays can be used to sequence multiple supplies.

Turning On Pass Elements

When using a "silver box" or "brick" power supply, turning on and off each voltage in a controlled order isn't always possible without additional circuitry. These power supplies provide standard voltages, such as 5 V, 3.3 V, 2.5 V and 1.8 V that are often distributed throughout a system.

For example, a brick can provide a 3.3-V logic supply and 1.8-V core supply to two different ICs. In some situations, these ICs will require different power sequencing; one device needs the core supply to rise first, while the second device requires that the I/O supply rises first.

One way to sequence supplies in this situation is to switch power through an external pass element. Fig. 4 shows a voltage detector connected to the gate of a MOSFET, which turns on and off V_{CC1}. An N-channel MOSFET is appropriate for this application when there is a higher voltage available to provide a gate-to-source voltage that's large enough to fully enhance the MOSFET.

However, a problem can occur during power-up of this circuit if V_{CC2} is present prior to V_{CC1} reaching a level sufficiently high to turn on the voltage detector's output. In this situation, V_{CC2} will enhance the MOSFET (i.e., it will be on) until V_{CC1} rises sufficiently for the voltage detector's output to assert a low.

This same sort of circuit can be realized with a voltage detector and a P-channel MOSFET without the need for a second, higher voltage. This circuit isn't suitable for low-voltage supplies, however. In addition, the higher on-resistance of a P-Channel MOSFET makes it impractical for high-power applications.

An easier and more reliable approach is to use a device, such as the MAX6819, to perform both the monitoring and the sequencing functions (Fig. 5). This type of IC monitors the first voltage with a reset circuit to determine whether it's within specification. When it is, the IC turns

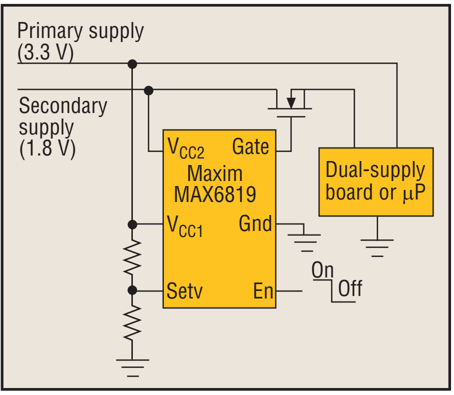

Fig. 5. *After the primary supply has powered up, the MAX6819 turns on the secondary supply. Its on-board charge pump enhances the MOSFET to minimize its on-resistance.*

on the MOSFET via its MOSFET driver. An internal charge pump adds a fixed voltage to the secondary supply and applies the resulting voltage to the gate of the MOSFET, which helps ensure the gate-to-source voltage is sufficiently high to fully enhance the MOSFET.

The Margining Function

During the manufacturing phase of many types of telecommunication, networking, storage and server equipment, a process called margining often is used to assess the robustness and future reliability of these systems. This technique involves an evaluation of the system (or processor), which is performed by deviating the power supplies from their nominal levels. To change those levels, it's common to adjust a dc-dc converter power supply by altering its feedback loop with a digital pot or a current DAC (Fig. 6).

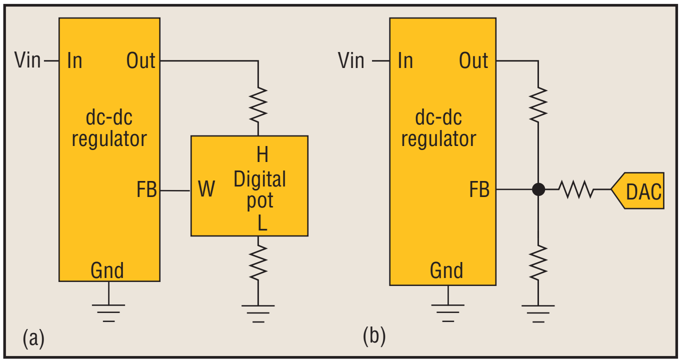

Fig. 6. *Two simple techniques for performing voltage margining include adding a digital pot (a) or a current DAC (b) to a dc-dc converter's feedback loop.*

Other common methods of margining include programming the supply's output through a digital interface or trimming the power supply. Different degrees of margining control include "pass/fail" approaches where you increase or decrease all power supplies to some level (e.g., ±5% or ±10%), or an finer-adjustment approach where you increment or decrement the supplies in smaller steps (e.g., 10 mV or 100 mV); the latter approach allows you to evaluate the system performance in greater detail.

An ADC can be used to measure these values more accurately. It's tempting to use the ADC contained within a

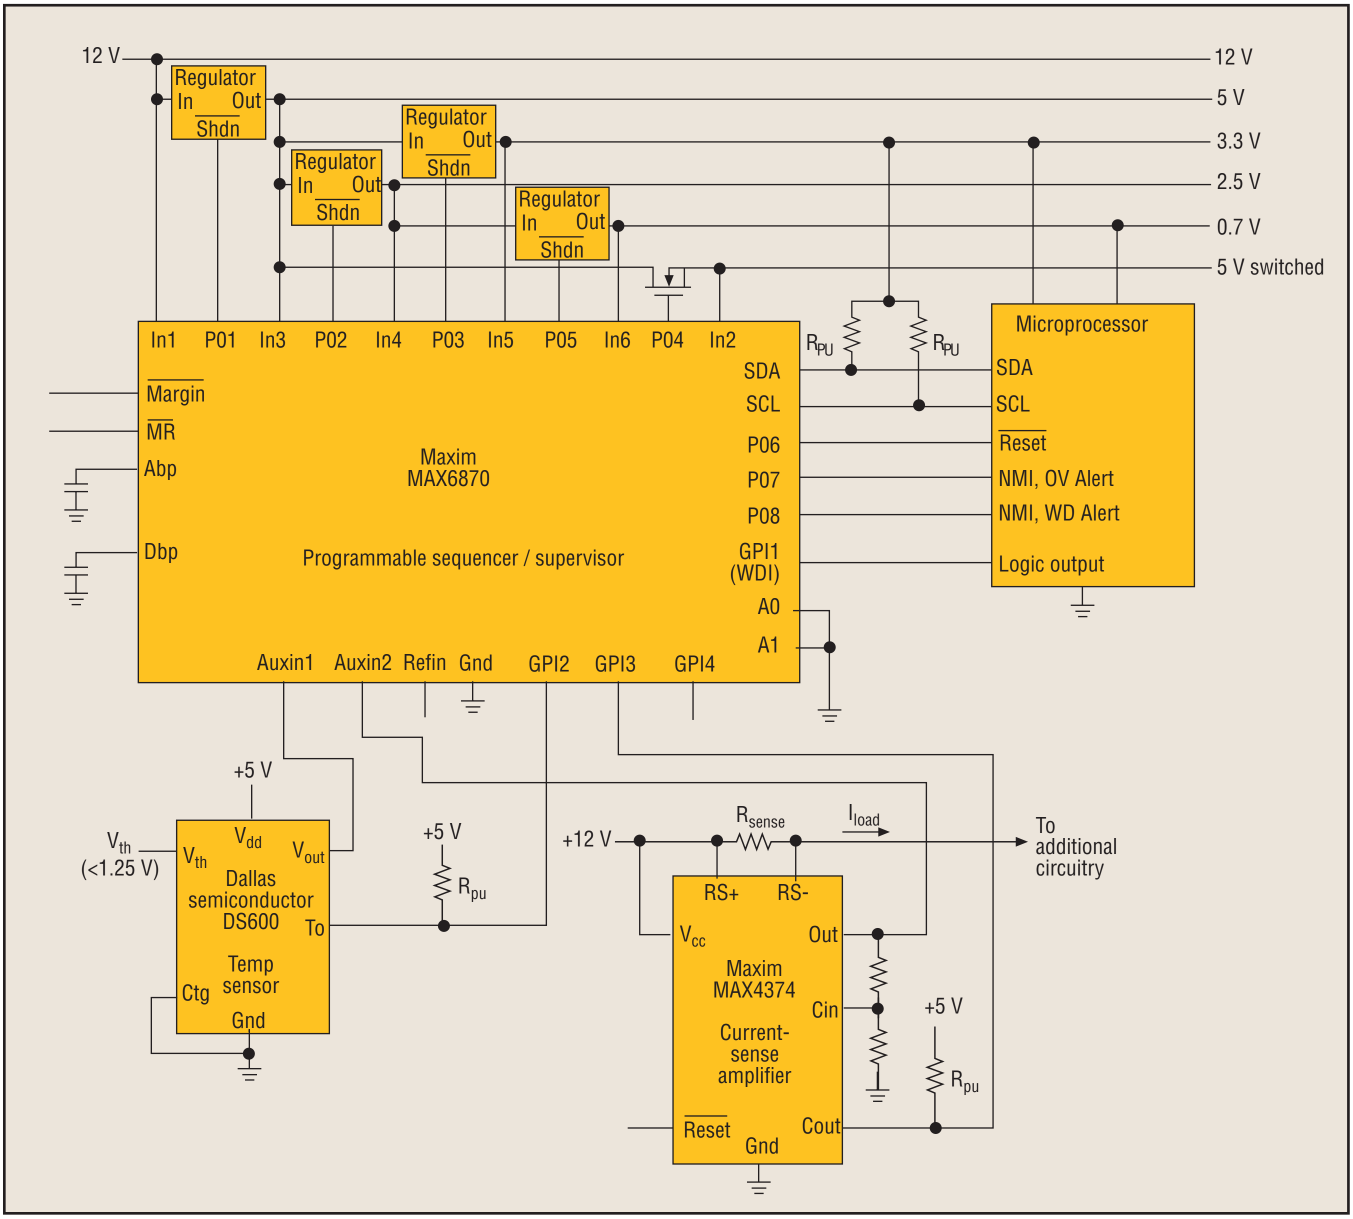

Fig. 7. *A programmable system-management device provides a flexible means for voltage monitoring and sequencing*.

microcontroller for this function; however, when the supplies powering the microcontroller drop below specification, its internal reference can be out of tolerance, thus affecting the ADC's accuracy.

Also, it's necessary during margining to disconnect or disable the reset output so the system can continue to operate. Otherwise, the system will reset, making it impossible to discover the supply-voltage levels at which the system fails. The process of performing these margining functions can be quite tedious when working with large systems.

Combining Monitoring, Sequencing and Margining

Although many processors require only two voltages—one to power the core and another to power the I/O—other devices, such as DSPs, ASICs, network processors and video processors, can require up to five supply voltages. Within a single system, supervisory circuits may need to monitor and sequence more than 10 voltages.

As the number of supply voltages in these systems continues to increase, the number of ICs needed to monitor, sequence and margin them also increases. Costs go up and more board space gets consumed. Also, when changes to parameters, such as a voltage threshold or a reset timeout period, are necessary, a new device may be required. Furthermore, changing the sequencing order becomes a difficult task.

One way to reduce the level of circuit complexity is to use a programmable system-management IC that combines the monitoring and sequencing functions. These new, fully integrated devices provide an alternative to existing solutions, giving more flexibility and function within a single device, thus saving board space, cost and design time.

The programmability of these devices makes changes easy to handle. They eliminate the need to swap parts in and out of a design during the prototyping and manufacturing stage. For many of these parts, a serial interface allows you to program the internal registers that both configure these devices and set threshold levels and delays; on-board EEPROM stores the contents of these registers.

Fig. 7 shows a MAX6870 system-management device set up to monitor and sequence several system supplies. When the +12-V bus voltage powers up and exceeds its threshold (stored within the MAX6870), one of the MAX6870 outputs enables the +5-V voltage regulator, either immediately or after a delay period (also stored within the device's memory).

After the +5-V regulator comes up and its output crosses its corresponding threshold, the +3.3-V supply is taken out of shutdown. The remaining supplies then power up in succession under this same scheme, except that the +5-V switched supply becomes available when the IC enhances the N-channel pass element.

You can usually program this type of system-management device to provide additional supervisory functions, such as reset circuits and watchdog timers. Furthermore, these devices can monitor parameters other than supply voltage through their analog and digital inputs. For example, in the circuit of Fig. 7, AUXIN_ (analog inputs) and GPI_ (digital inputs) monitor a temperature reading and a power-supply current-sense reading.

The MAX6870 includes a 10-bit ADC that digitizes those readings, and the microcontroller monitors the status of the digitized readings. The temperature sensor and the current-sense monitor each include a comparator output that indicates a fault has occurred (i.e., the temperature or the current has exceeded a specific limit). Each comparator output connects to a general-purpose input (GPI) within the system-management device. In addition, the IC can be configured to turn off one or more power supplies when these fault conditions occur, thus lightening the load on the +12-V supply.

The internal ADC makes it easy to accurately margin a power supply. The voltage at each power supply's output can be read from the ADC registers during the margining process. Also, the margin input can disable the outputs or program them to a known state when the supplies are margined, preventing the system from resetting during this procedure. **PETech**

For more information on this article, CIRCLE 334 on Reader Service Card