Electronic Design

By FREDERIK DOSTAL, Analog Devices

Power Sequencing, Simplified

A new approach to power sequencing doesn't require any programming, digital communication, or even a bias power supply.

From graphics processor units (GPUs) and central processing units (CPUs) to field-programmable gate arrays (FPGAs), digital signal processors (DSPs), and heterogeneous SoCs, every processor requires power sequencing. These digital loads often operate from multiple voltage rails, each of which may supply a different on-chip power domain. Power sequencing is the process of applying and removing power from these voltage rails in a specific order during system startup and shutdown.

If the appropriate sequencing isn't followed, the function of the load may be limited and defective. Also, in extreme cases, the load can be destroyed. To prevent this, sequencing functions are employed.

This article discusses different possibilities for the implementation of a sequencing and voltage-monitoring function. It also introduces a new, simple-to-use sequencing solution that requires no programming, digital communication, or bias power supply. These features are important because the sequencer itself needs to be powered on first in a system. By that time, processors or other loads usually aren't running yet.

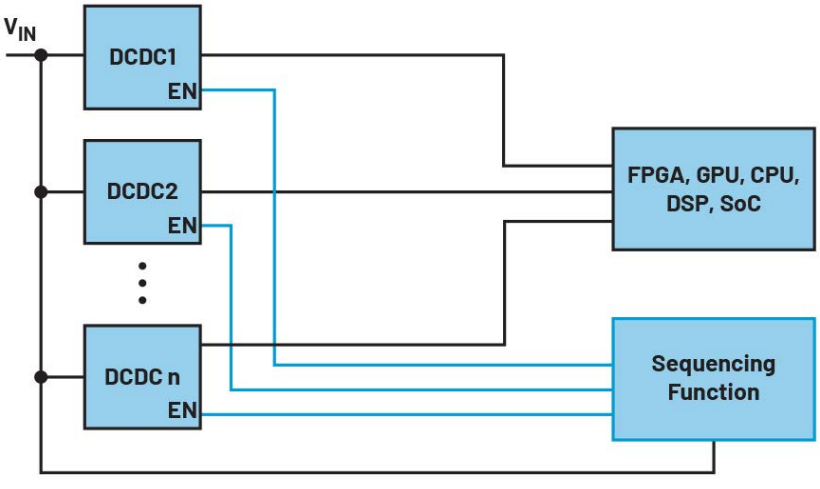

1. DC-DC converters supplying loads with multiple voltage rails need to coordinate their power-up and power-down sequences to ensure proper system operation. Analog Devices

How Power Sequencing Typically Works

One common approach to sequencing involves placing several DC-DC converters in front of the processor and switching them on one after the other by linking the generated voltages with enable (EN) pins. One drawback with this approach is the requirement of voltage regulators with an exact definition of the precision enable pin. The other drawback is that a fixed sequence of switch-off of the individual voltages is difficult to implement. Also, the exact time course can only be implemented with additional cumbersome circuitry.

Given all that, a dedicated sequencing controller is often preferred. *Figure 1* outlines such a system. The voltage converters — from the first, DCDC1, to the last, DCDC n — are switched on and off by a sequencer in the correct order through the respective enable pins.

Ideally, the sequencing circuit should support a wide input voltage range since it needs to perform its function before the respective voltages in the system are switched on. Thus, a wide operating input voltage range is favorable.

The sequencing function can be implemented in several different ways. A microcontroller (MCU), a small FPGA, a dedicated programmable power sequencer, or a simple sequencer could be used, depending on the complexity and flexibility required.

Many of these solutions have unfavorable limitations. The MCU must first power up and execute its firmware be-

fore it's able to perform the sequencing function. A small control FPGA or a programmable power sequencer must be programmed before operation. This can be completed before the device is soldered onto the PCB or after assembly, prior to the system's initial power-up. In both cases, the additional programming step can lead to considerable effort and costs.

A New, Simple-to-Use Approach to Power Sequencing

An alternative approach eliminates the need for programming in production. In this type of sequencing solution, timing and operating behavior are configured using only external resistors and capacitors. With a very wide input voltage range of 2.7 to 15 V, this sequencer can be operated directly from a 3.3-, 5-, or even 12-V supply voltage without

SYSTEM DESIGNERS must contend with the increasing complexity of power sequencing. Today's electronic systems need more power density, with the number of voltage rails growing at both the chip and system levels.

Today's processors are more heterogeneous than ever. Everything from CPU cores are crowding onto the same silicon die as on-chip accelerators, such as GPUs and NPUs, as well as memory controllers, high-speed SerDes, and I/O interfaces like PCIe. Each of these subsystems may require a different supply voltage to run efficiently.

To manage this complexity, chip designers partition the processor into independent power domains, often with dedicated voltage rails. This enables fine-grained power management, allowing idle blocks to be completely shut off or dialed down independently to save power.

As a result, a single processor today may contain a dozen or more voltage rails and power domains, while the complete system can include dozens or even hundreds of interdependent power rails and voltage domains.

Supplying power to all of these rails without considering the power-up and power-down timing for each one poses several dangers. Improper sequencing can cause system failures while increasing the likelihood of long-term damage to the main processor or any of the chips surrounding it, e.g., FPGAs, DSPs, and analog-to-digital converters (ADCs), among many others. One of the primary concerns is the possibility of breakdown in the electrostatic-discharge (ESD) protection and isolation barriers that internally separate two power rails in a chip.

If one rail remains powered while another is inactive for extended periods — or repeatedly experiences these conditions over time — such protections can break down, potentially causing permanent damage to the chip.

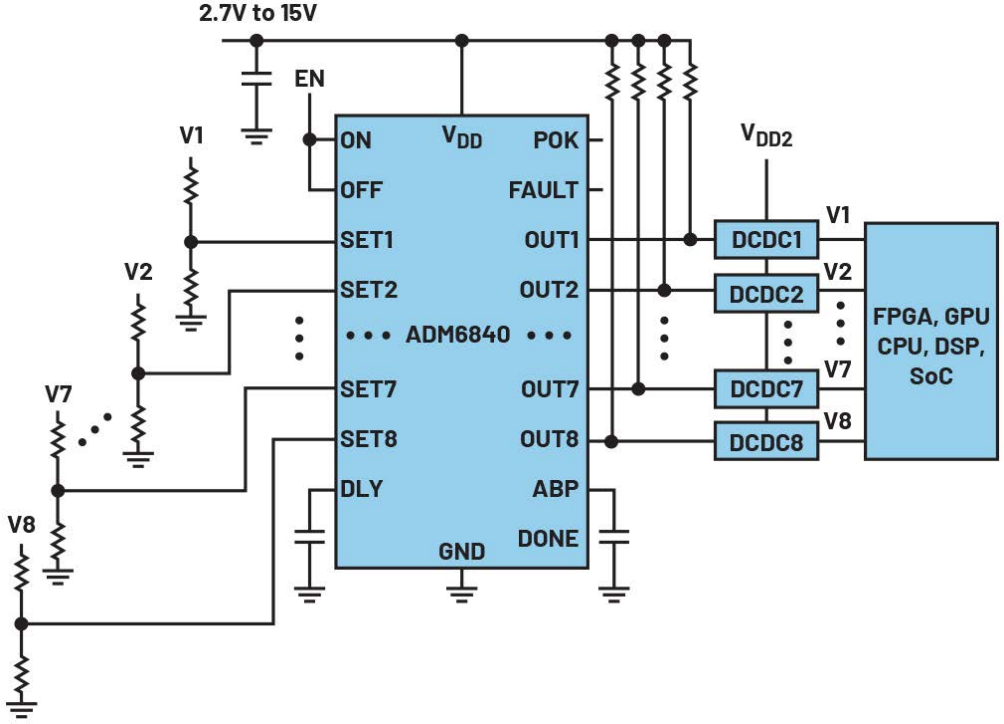

2. The ADM6840 requires no programming and supports sequencing and monitoring of up to eight voltage rails. Analog Devices

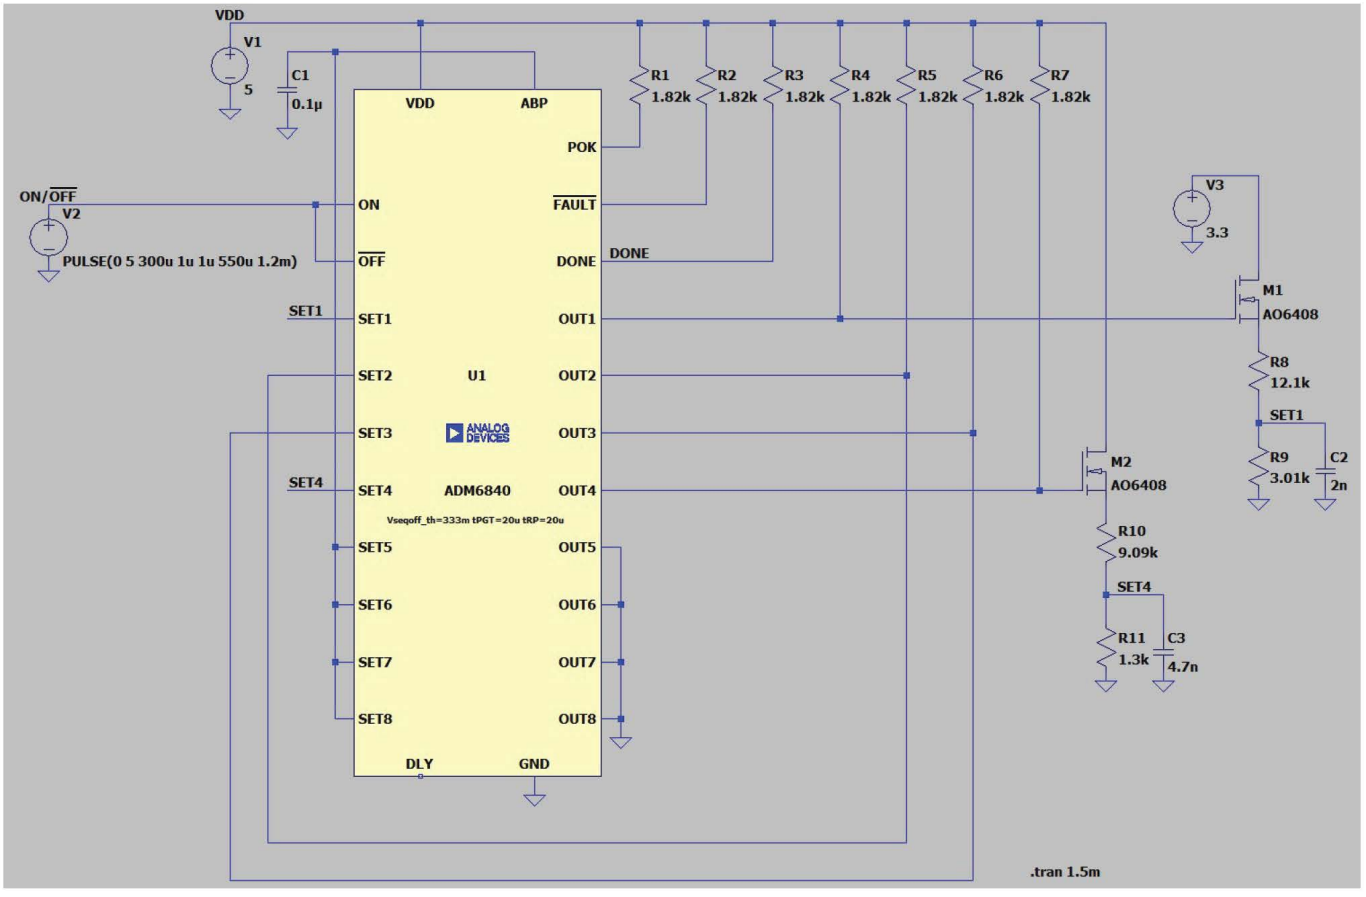

3. LTspice can be used to evaluate the ADM6840 power sequencer and supervisory device. Analog Devices

needing an additional power supply for the sequencer itself.

Another advantage, as seen in *Figure 2*, is that the different voltages (from V1 to V8) are monitored separately. Only when all voltages have reached a predetermined value, a power OK signal is given. One of the most important metrics when it comes to voltage monitoring is accuracy. Monitoring accuracy is a part of the overall accuracy of each voltage generated by a DC-DC converter. The ADM6840 comes with a voltage monitoring accuracy of ±0.8%.

While accuracy is a critical factor for the power sequencer, reliability is also key to keeping the circuit well-protected. The ADM6840 can be operated in a temperature range between –40 and +125°C, which is helpful since many circuits often become quite hot. If more than eight channels are needed, several sequencers can be operated together.

Evaluating the function of the ADM6840 sequencer can be simplified by using LTspice. Within LTspice, the sequencer is available in an example circuit. *Figure 3* shows the simulation environment in which different setups can be simulated.

Selecting the Right Components for Power Sequencing

Powering a modern, multi-rail processor is challenging enough without adding the complexity of power sequencing. With the new solution, however, sequencing can be straightforward to implement. In most applications, no additional auxiliary power supply is required, nor is there any need for complex, time-consuming programming. As a result, engineers are able to save themselves development time, while ensuring safe, reliable system startup and shutdown.