

Advanced System Management with Analog Non-Volatile FPGAs

Instant-on Intel® MAX® 10 FPGAs with integrated ADCs provide crucial system management device features.

Author Introduction

Rich Howell

Sr. Product Marketing Manager
Intel Programmable Solutions Group

A system management device is critical to the successful power up, configuration, maintenance, and power down of the system. The requirements of these devices are growing as system complexity increases. System management device features such as instant-on, analog capability, and flexibility are crucial. Instant-on Intel® MAX® 10 FPGAs with integrated analog-to-digital converters (ADCs) offer fast start-up time and are capable of monitoring many voltage rails and/or temperature sensors. Embedded flash for data logging, combined with support for soft Nios® II microprocessor provides a robust diagnostics and prognostics platform. MAX 10 FPGAs enable the flexibility required to customize the system manager to specific board requirements. The combination of these attributes makes MAX 10 FPGAs an ideal system management solution.

Industry challenges

System management is used to describe the many tasks and functions required to start-up, configure, maintain, and power-down a system, while minimizing costs and maximizing utility and reliability. These tasks are common across a wide variety of application boards especially where large datapath FPGAs and ASSPs are utilized.

As process nodes shrink devices, datapath FPGAs can integrate more functionality into a single device. For example, Intel Arria® 10 FPGAs integrate up to 1.15 M logic elements (LEs) and 96 transceivers running up to 28.1 gigabytes (GBs) for an unprecedented level of integration. At the same time, to support these new capabilities, an Arria 10 FPGA has nine power rails that must be sequenced and monitored to maximize system performance and operational uptime.

The application boards used for system management represent a significant investment by both developer and system provider. To protect this investment, monitoring and controlling board-level environmental conditions is critical. Operating a system out of recommended conditions can not only affect performance, but can reduce system reliability, life expectancy, and even lead to catastrophic board failure.

The service providers who use these complex boards need to maximize their investment and, ultimately, operational up time. A system must be able to run a self-diagnostic test and determine that the board is operational. A good board will be able to indicate the nature of any failure so that any service call can be very targeted and efficient to minimize down time. An intelligent system will be able to predict when a failure is imminent and alert the service provider before the failure, and eliminate system down time altogether.

Table of Contents

- Introduction 1
- Industry challenges..... 1
- Power rail management..... 2
- Thermal management 3
- Diagnostics and prognostics 3
- Available today..... 4
- Where to get more information... 4

A robust system management design incorporates a wide variety of tasks in both the analog and digital domain including power rail management, environmental condition management, and analytics for diagnostics and prognostics. An instant-on mixed-signal FPGA, like Intel MAX 10 FPGAs, offers the right mix of analog and digital resources in a single chip.

Power rail management

Power rail management is commonly recognized as one of the most critical functions in system management, and has become increasingly important with smaller process geometries. The power state of a board is not a simple binary on/off. The proper power control of a board is more involved than simply turning 'on' the power. As seen in Table 1, there are several different power states in which a board can reside, each with its own requirements.

Today's leading-edge 20 nm datapath solutions, like Arria 10 FPGAs, have multiple power rails, at different voltage levels, each with specific tolerances and sequencing requirements. Figure 1 shows the power up and monitoring of a typical Arria 10 FPGA design using MAX 10 FPGAs with integrated analog capability as the system manager. MAX 10 FPGAs contain one or two integrated 12 bit successive approximation register (SAR) analog-to-digital (ADC) converters. There are 16 analog inputs multiplexed to the ADCs enabling the MAX 10 FPGA to monitor a large number of analog signals.

An instant-on system management device, the MAX 10 FPGA can control and initialize the rest of the board. Using an analog input, the MAX 10 FPGA will monitor the line

side of the primary power rail (A0) and if within operating tolerances, power-on state, enter the power-up stage. Turning on the Intel® Enpirion® Power Solutions using digital logic and general-purpose I/O (GPIO) (D1 – D6) in a controlled manner, the MAX 10 device ensures the Arria 10 devices power up correctly. By monitoring the line side of the power, the MAX 10 devices ensure that each rail is operating at the correct level (A0 – A6), and then it can release the Arria 10 FPGA from reset for normal operation.

The MAX 10 FPGA, with up to 16 analog inputs, has sufficient analog resources to monitor the multiple voltage rails for the Arria 10 FPGA. After the Arria 10 FPGA is released from reset and is operating, the MAX 10 FPGA can continue to monitor the system power rails. A faulty power supply is not always easy to diagnose from a system level. If a power rail 'dips' out of tolerance, it can put the device (or parts of the device) into an unknown state. In reality, this means device operation cannot be guaranteed. Portions of the device may continue to operate correctly, but other parts may not. So it is important for the system manager to stay vigilant in monitoring the power rails and if there is a deviation, it can reinitialize the system as needed.

Power consumption continues to become an increasingly important aspect to datapath designs for a number of reasons, including local thermal management cost of operation, and facility-wide cooling requirements. To help offset facility cooling, it can be advantageous to power down some or all of the board FPGAs during times of reduced work load. In our previous example, we show a single Arria 10 FPGA, but in systems where multiple Arria 10 FPGAs are

Power State	Description	Key Design Consideration
Off	No power applied at any level.	Some boards require hot-socketing functionality. In these cases, the system manager must be able to support hot-socketing to initialize the board upon insertion.
Power on	Main power supply is on and stable, but board power tree is not yet initialized.	If line side power is not at nominal operating levels, this could be a sign of an unstable power supply leading to brownout conditions and preventing a stable system initialization.
Power up	Board power rails are initialized at the prescribed ramp rate and sequence to ensure stable board operation. This is when FPGAs are configured and microcontroller units (MCUs) are booted.	Power rails that are not sequenced properly can lead to excessive power consumption, unstable device operation, reduced device operating life, and even device damage.
Run time	Normal system operation. Power rails must maintain proper voltage levels and provide sufficient current for proper system operation. System health and status is communicated to the higher level system.	Power rails that drift over time, or fluctuate beyond nominal ranges, can put devices into unknown or unstable states. If the variance is not detected, unstable operating conditions could persist, reducing system reliability.
Low power / sleep (optional)	In complex systems with multiple datapath FPGAs or with multiple blades in a chassis, during periods of lower activity, the ability to turn down or turn off portions of the system can greatly reduce power consumption, cooling requirements, and operating cost.	Putting a device into a power down or sleep mode, often has a specific voltage control and sequence depending on whether the device is being powered down completely, going into sleep mode, maintaining volatile state conditions. These requirements will vary by device.
Power down	Powering down the system power rails in the required sequence. The reverse of power up.	Uncontrolled power down of a device's power rails could create internal potential differences that can reduce a device's operational life.

Table 1. Possible Power States of a System Board




Figure 1. MAX 10 FPGA Power Sequencing of Arria 10 FPGA

employed, the MAX 10 FPGA can control them independently and power one or more of them as application needs dictate.

Large datapath FPGAs will sometimes have a specific power-down sequence to ensure long-term reliability of the device. Just as the MAX 10 FPGA initialized the system at startup, it has the right mix of analog and digital resources to power down the Enpirion Power Solutions in the required sequence to power down the system in a safe and controlled manner.

Thermal management

Thermal management is considered to be the next most critical aspect of the system management design. A system that does not operate within the appropriate thermal range of its components will be more prone to reliability issues, early wear out, and potentially catastrophic failure. The system manager must be able to monitor the environmental conditions and evaluate and analyze the results of the data gathered. Ultimately, the system manager must be able to action as needed to mitigate poor conditions and or initiate steps to prevent system damage.

In complex systems with multiple datapath devices, it is possible for board-level micro climates to exist. One datapath FPGA may be more heavily loaded than another; therefore, it is good design practice to track the temperature at multiple locations on the board. This ensures complete coverage of board conditions. The system management device must have sufficient resources, analog and digital, to monitor multiple temperature sensors and the ability to control active cooling systems, often fans. There are several different types of temperature sensors, the most common interfaces include analog, I²C, and serial peripheral interface (SPI). The MAX 10 FPGA with integrated ADCs enables it to connect to analog sensors, and as an FPGA it can easily support both the I²C and SPI interfaces, or all three variants in the same device.

Diagnostics and prognostics

Even boards with robust system management implementation will sometimes fail. Even though the system management design was not able to prevent board failure, its job is not done. The system management device should be able to record events prior to, during and after the failure.

Let's reuse our Arria 10 FPGA design example from earlier. In this case, there is a temperature sensor next to the Arria 10 device and it shows an increase in temperature. As long as the temperature is still within the normal operating range, the system manager would increase active cooling (for example, increase fan speed). The temperature continues to increase and crosses a threshold approaching critical. In this case the system may choose to do any of the following or a combination of them: increase active cooling, increase the frequency of monitoring, log the event and possibly event redirect traffic reducing the loading on the device in question. Ultimately if the temperature continues to rise to a critical threshold, then the system management system should shut down the device or entire board as necessary.

In the event of a system shutdown due to failure, the system log that captured the various environmental excursions becomes very important. This should be transmitted back to the system provider to expedite and guide system replacement and maintenance. It can also be used by the system provider and equipment OEM to identify any design and/or environmental improvements.

The data log is a great aid for system repair as well as for conducting a post mortem on a system. However when we look beyond diagnostics to prognostics, data logging can become more important by predicting failures before they happen. A power rail that is dropping down towards a threshold limit, but hasn't reached it yet, could be close to failure. An FPGA that continually 'runs hot' could indicate a cooling system that is starting to lose effectiveness or is subject to a foreign object obstructing air flow. It is better to know of an impending failure before it happens. That way preventive measures can be taken to maximize overall system up time. Either network traffic can be diverted around the hampered system or repairs can be made before the system goes off line. With large blocks of embedded flash for data logging and support for a Nios II soft processor, the MAX 10 FPGA is well suited to capture system fault conditions for diagnostics as well as analyze behavior for possible failures that have not yet occurred.

Available today

MAX 10 FPGAs are available for purchase today. Visit <https://www.altera.com/max10> to learn more.

Where to get more information

For more information about Intel and MAX 10 FPGAs, visit <https://www.altera.com/products/fpga/max-series/max-10/overview.html>

