System Management Using a Mixed-Signal FPGA

December 2006

Table of Contents

List of Figures

List of Tables

System Management Trends and Challenges

System management continues to gain importance in the design of all electronic systems because smaller process geometries drive more multi-volt devices and are more susceptible to voltage and temperature fluctuations. System management is a collection of seemingly unrelated tasks with the goal of ensuring the proper operation of the system. These tasks focus on maximizing system uptime, identifying and communicating alert conditions, and logging data and alarm conditions. Boards must also be able to log an event, initiate corrective action, and/or notify a remote supervisor when fault conditions occur. Driven by the need to increase system uptime and reliability, many systems are adding in-system diagnostics and prognostics, not only to help debug systems that have failed, but also to identify potential failures before they arise. In markets driven by standards, reliability and uptime are key metrics by which OEMs can differentiate themselves.

System management tasks often serve varying roles at different operational stages of the board/system, as shown in Table 1. In this case, six stages of system operation have been identified: Power-On, Power-Up, System Operational, Sleep / Low Power, Power-Down, and Off. Table 1 lists the various system management functions and their relevance to the various operational stages.

- Power-On is defined quite simply by the primary power supply to the board being up and stable. Once this supply is stable, the board can begin to power up.
- Power-Up is the period in which all system power rails are initialized. System clock signals are initiated and stabilized. MCUs and volatile FPGAs are booted, configured, and/or initialized. Local and remote communication channels are established. Reset can then be released for board operation.
- System Operational defines normal operational mode.
- Sleep / Low Power Mode is an optional mode that can be utilized by some applications where 100 percent operation is not required. The system may elect to turn off portions of the board to conserve power, waking up at predetermined times and/or upon external stimuli.
- Power-Down involves shutting down the system in a controlled manner to ensure no damage is done to components on the board. Also, many systems may need to save data and the "state" of the system so the next power-up can be implemented properly.
- Off indicates when no power is supplied to the board. No system management tasks are active, but the potential for hot-swap or plugging into a live system exists, so support for this capability is required in many designs.

Table 1: **System Management Functions**

	Power-On	Power-Up	System Operational	Sleep / Low Power Mode	Power-Down
Power-On Detection and Reset	✓				
Power Sequencing		✓			✓
Reset Management	✓			✓	
Voltage Monitoring	✓	✓	✓		✓
System Clocking	✓	✓	✓	✓	✓
Data Logging	✓	✓	✓		✓
Remote Communications	✓	✓	✓		✓
Diagnostic and Prognostics	✓	✓	✓		✓
Error/Alarm	✓	✓	✓		✓
Current Monitoring		✓	✓		
SRAM FPGA Management		✓	✓		✓
Thermal Management			✓		
MCU Boot Loader			✓		
ID and Authentication			✓		

Today's Solution

Current system management implementations require a large number of discrete components (sometimes numbering in the hundreds), occupy large amounts of board space, and are inflexible to change. These solutions are a collection of fixed-function chips (Figure 1) and discrete components that must work in concert to create a cohesive solution: CPLD, real-time clock, power sequencer, temperature monitor, fan controller, nonvolatile memory, PWM, configuration memory, etc. In addition to consuming board space, the large number of components adds to cost both directly (unit cost, assembly cost, and inventory cost) and indirectly (design time, procurement, and discontinuation). Increased component count also contributes directly to the failure rate; eliminated parts will not fail. These hardware-implemented discrete solutions often require component changes or board respins, even for incremental design changes, making it impossible to create platform solutions and, more importantly, requiring requalification.

Most system management implementations are proprietary, having evolved over time within an organization, although standards are being developed and adopted (ATCA, µTCA, and IPMI).

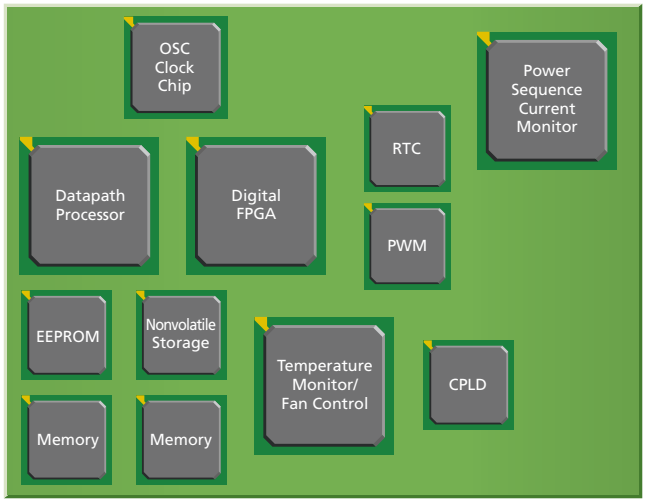

Figure 1: **Typical Board Using Current System Management Solution**

System Management in a Single Chip

As the world's first mixed-signal FPGA family, Fusion devices integrate mixed-signal analog, Flash memory, and FPGA fabric in a monolithic device. Actel Fusion™ devices enable designers to quickly move from concept to completed design and deliver feature-rich systems to market. This new technology takes advantage of the unique properties of Actel Flash-based FPGAs, including a high-isolation, triple-well process and the ability to support high-voltage transistors to meet the demanding requirements of mixed-signal system design.

The Fusion family sets the standard in system management as the market's only single-chip solution. Leveraging a unique combination of Flash memory, configurable analog, and Flash FPGA logic, Fusion architecture reduces cost, board space, and design time.

The flexibility of Fusion architecture enables designers to implement proprietary or standards-based system management solutions. As a single-chip implementation, a Fusion device simplifies design implementation, enabling all system management functions to be configured in a single design environment and increasing reliability.

Figure 3 shows a system that uses the Fusion mixed-signal PSC in a system management application. Many of the elements from Figure 2 are incorporated into the Fusion device, including the control-plane microprocessor, such as CoreMP7 or Core8051, and Flash memory.

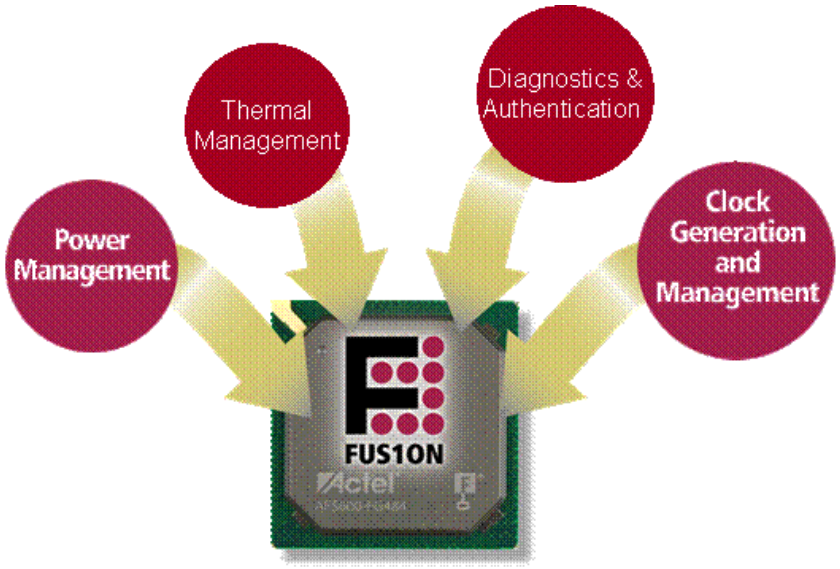

Figure 2: **Fusion Single-Chip Solution**

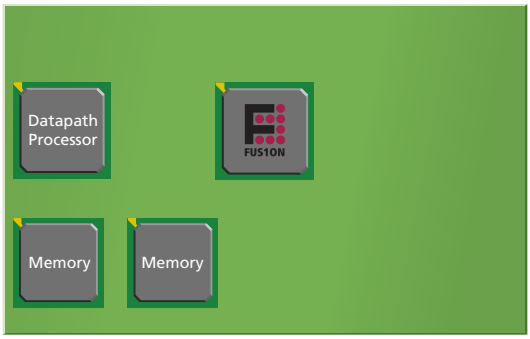

Figure 3: **Typical Board Using Fusion System Management Solution**

System Management Tasks

As stated previously, system management involves a variety of tasks. A number of the critical tasks and the benefits of Fusion devices for each task are discussed in this section of the document.

Power Management

As the most recognizable aspect of system management, Fusion architecture takes aim at power management with an I/O structure specifically designed to effectively and efficiently manage multiple power rails directly (refer to Figure 4). Fusion can manage up to 10 power supplies and connect directly to power rails from –11 V up to 12 V without any external components. The Fusion Analog Quad (AQ) I/O structure combines three analog inputs with a gate driver output. Each analog I/O is individually configurable and can operate independently. Additionally, these I/Os were designed and grouped together to work in concert to support power management. The robust functionality of the AQ allows for the following:

- Power-on detection and reset
- Power sequencing
- Voltage monitoring and trimming
- Current monitoring

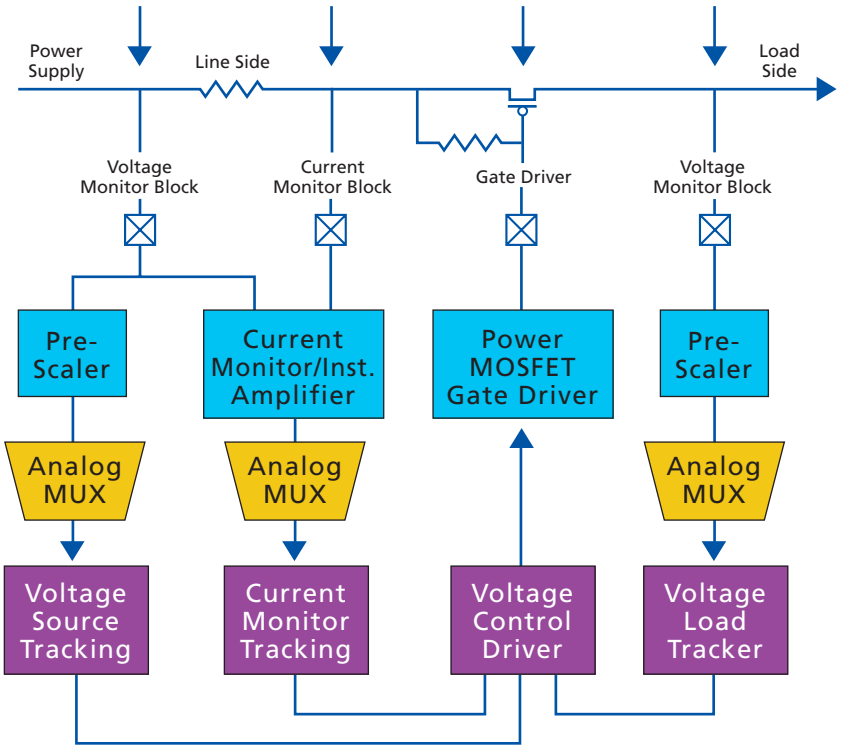

Figure 4: **Voltage Rail Control Example**

Power-On Detection and Reset

The Fusion PSC is live at power-up and can operate from a single 3.3 V supply with no power sequencing requirements of its own, making it an ideal system master. It can monitor various power supply rails, keeping them gated from the rest of the system until they have stabilized. Once the supplies have stabilized, the programmable gate drivers can turn on the power rails in the appropriate sequence and at

the appropriate ramp rate. A designer determines the sequence by looking at the needs of the components used and the power supply load balancing. During the power sequencing process, Fusion devices can hold the system and its various components in reset, ensuring stable system initialization.

Power-Up Sequencing and Tracking

Fusion devices can support a number of different methods for power rail sequencing. The least sophisticated methodology would be to turn on all the rails simultaneously. However, there are some drawbacks to this strategy. Many devices, such as SRAM-based FPGAs, suffer from inrush current at power-up, often over 1 A per device. If there are multiple such devices within a system, having them turn on simultaneously will require a larger, more expensive power supply than is needed for system operation, just to accommodate this inrush. Bringing up these devices one after another will better balance the system and enable the designer to properly size the power supply. Utilizing the on-chip real-time counter (RTC), a time-base sequence can be implemented. Ramp up one voltage rail, wait a predetermined amount of time, then ramp up the next voltage rail.

The on-chip ADC can be used to monitor the load side of the rail during power-up so that this can be used as a trip point for initiating the next rail in the sequence. Ramp up one voltage rail, then when it reaches a certain level, ramp the next rail in the system.

Some multi-volt devices require that their rails come up together such that there is no more than a diode drop between them, which could forward-bias structures within the chip. The AG pads used to gate individual power rails are configurable, allowing for finer control over the individual ramp-rates. Fusion architecture vastly reduces the number of external components required to monitor and sequence power supplies, reducing board space, component count, and cost.

Voltage Monitoring and Trimming

To ensure proper system operation, voltage rails must stay within their defined limits. The Fusion device can monitor power rails during operation for both brownout and over-voltage conditions. Additionally, when combined with CorePWM, I/Os on the device can be used to trim voltage levels of trim-capable power supplies. This has the benefit of proactive control over voltage levels and enables the use of lower cost power supplies (with wider tolerance specifications) because they can be actively tuned.

Current Monitoring

Monitoring total power consumption enables designers to better plan for thermal conditions and system loading. Unexpected increases in current levels can point to potential failures or other alarm conditions. For example, in large industrial motors, a rise in the current to the windings often indicates a need for service or replacement. The Fusion architecture's integrated current monitoring block uses a differential input to monitor the voltage drop across a small, user-selected external sense resistor.

The Fusion device vastly reduces the number of external components required to monitor and sequence power supplies, reducing board space, component count, and cost.

Thermal Management

Maintaining the proper environmental operating conditions is a key component to system management. Today's intelligent systems not only monitor and manage thermal conditions, but also distribute system traffic to better balance the system and maximize performance. Fusion devices integrate temperature monitoring capabilities into the Analog Quad, requiring only an external bipolar transistor. A Fusion device can easily monitor up to 10 remote temperatures, so designers can track the temperature of not only power converters and inlet or outlet air, but also power-hungry FPGAs and processors.

In addition to temperature monitoring, Fusion devices can handle fan control to enable closed-loop thermal management (Figure 5). The Fusion architecture's integrated thermal management enables designers to easily and effectively maintain optimal system conditions, leading to increased uptime and performance with fewer components and reduced cost.

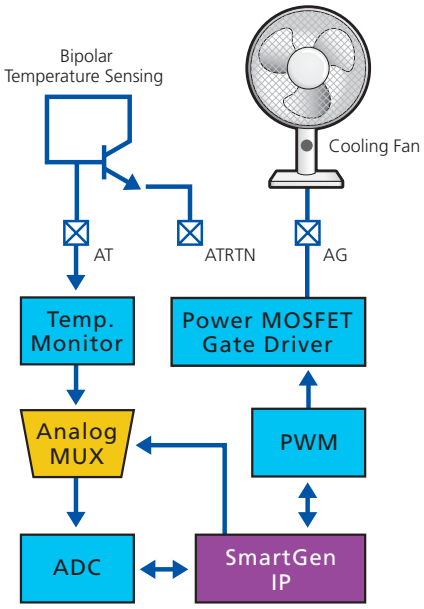

Figure 5: **Fusion Thermal Management**

SRAM FPGA Management

FPGAs continue to outpace market growth, replacing ASICs in many applications. A drawback to SRAM FPGAs is the hand-holding required to keep them operational. At power-on, volatile SRAM devices must be configured with their designs and often require voltage sequencing to avoid programming errors or other problems on-chip. Fusion integrated Flash memory can store the design files for many different types of FPGAs, eliminating a separate configuration PROM from the board. Since SRAM FPGAs are volatile devices, the user must ensure their configuration memory is not upset by voltage dips. This can be avoided by using a nonvolatile FPGA technology such as Flash or antifuse, but for those do use SRAM technology, brownout detection is another additional burden associated with using SRAM. With many SRAM-based FPGAs, a voltage variation as little as 60 mV is enough to put the device in an unknown state requiring an FPGA reset.

The Fusion architecture's unique feature set makes it an ideal SRAM FPGA management device, eliminating boot PROMs and brown-out detection devices, further reducing board space and overall cost.

System Clocking

Proper clock generation and distribution is critical to system operation. Whether the clock source is external to the system or internally generated, the Fusion architecture's robust clock distribution and generation resources are up to the task. Fusion devices support both internal and external clock sources, including an integrated 100 MHz RC oscillator requiring no components (Figure 6 on page 9). The 100 MHz RC is accurate to 3 percent over the industrial temperature range. For systems that require higher precision, an external crystal can be connected to the Fusion crystal oscillator circuit.

Integrated clock conditioning circuits (CCCs) and phase-locked loops (PLLs), combined with high-speed, low-skew networks, allow for the manipulation and distribution of very accurate clock sources to both on-chip and board-level components. Clock sources available to the CCCs/PLLs include the internal RC and external crystal oscillators, an external clock source, and an internal signal. For systems dependent upon an external clock source, the RC oscillator serves as an excellent backup clock, ensuring your board never gets hung out to dry in the event of primary clock failure. Fusion devices increase system reliability by reducing clocking chip count and providing an integrated secondary clock source.

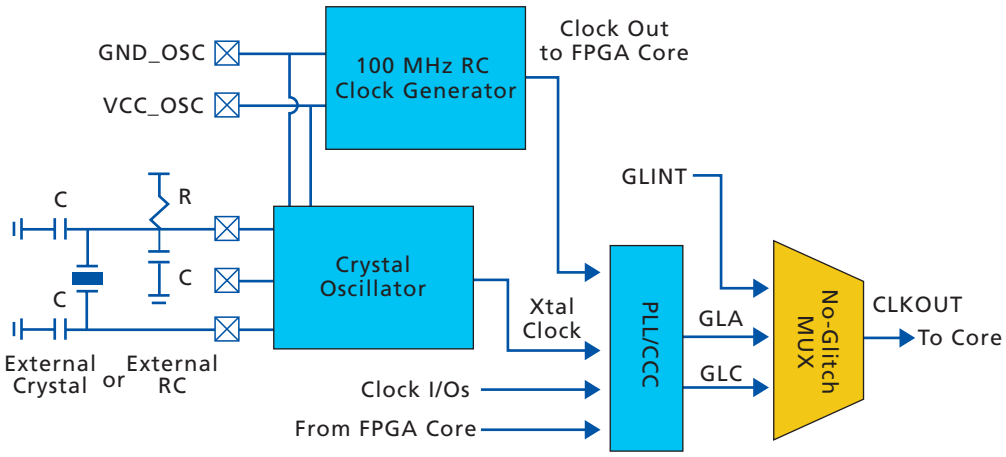

Figure 6: **Fusion Clocking System**

The Fusion RTC includes a programmable 40-bit counter and match register (timer) that generates time-based match events. In addition to use models such as watchdog timer, device lifetime monitoring, and event timer, the RTC can be used to wake the Fusion device out of Standby mode, enabling very low power modes of operation. The No-Glitch MUX (NGMUX) allows the device to switch between asynchronous clock domains in a controlled manner. The device switches to a slower clock frequency during periods of relative inactivity, thereby reducing active power consumption.

Diagnostics/Prognostics

The ability to determine failure modes and even predict them is quickly gaining momentum in system management support. The ability to read back timestamped system parameters about board operation is invaluable to failure analysis. The on-chip Flash memory of Fusion devices provides designers the ability to save and timestamp key system parameters, such as current consumption of power rails, device temperatures, and voltage rail fluctuations. Not only can this data be analyzed post-failure to identify the root cause of failure, but innovative designers are looking to analyze system trends during operation. By analyzing how a particular parameter varies over the life of the board, it is possible to predict a failure before it occurs, thereby increasing system uptime. Using the industrial motor control example again, a customer might measure the current to the windings as well as motor vibration (as a voltage input) to determine when to bring the equipment down in a planned fashion. In industrial applications, a planned shutdown is dramatically less expensive than an unplanned one due to the cost to fix the problem as well as the possibility of lost profit from equipment shutdown. Fusion architecture enables you to send in the repairman before the board fails!

Remote Communications

As a programmable device, Fusion can support a wide variety of communication protocols and implementations for both local and remote communication standards. There are a number of Actel DirectCore and CompanionCore IP blocks available for standards-based communication, and users are free to implement their own proprietary communication protocols.

Communications that can be implemented with Fusion include, but are not limited to the following:

- GPIO—Large number with many I/O standards supported
- UART, I2C (IPMB-0 & IPMB-L), SPI, SMBus, PCI, proprietary, etc.
- Slave and master modes
- Ethernet, proprietary

Identification/Authentication

Security continues to be a primary concern for users and developers alike. Fusion devices are designed with security in mind. As a single-chip solution, there is no communication between devices of any programming files that could be intercepted and cloned. Backed with AES security, Fusion devices are secure from external tampering. Designers can implement secure handshakes or device authentication to ensure only authorized components are used within a system. This prevents the manufacturing of knock-off system cards or other devices that can affect system reliability, your reputation, or your bottom line.

Tools/Solutions

The Actel System Management Development Kit provides an excellent platform for developing system management applications or applications with a microprocessor. The kit includes an ARM®-enabled Fusion device and a System Management GUI, and provides a platform for systems that perform the following functions:

- Power-up detection
- Power sequencing
- Thermal management
- Sleep modes
- System diagnostics
- Remote communications
- Clock generation and management

The board includes an M7AFS600 device, which is compatible with CoreMP7, Core8051, and other processors as well as non-processor-based implementations. The kit's demonstration design uses the MicroC-OS-II on CoreMP7 with the System Management GUI, which is multi-tabbed for board status, application data, IPMI data, and a graphical display of the monitored analog data. The board status tab enables the user to monitor all of the power supplies, set the LEDs, set and read the RTC, read on-board and on-chip temperature sensors, and display the contents of the embedded Flash, including event logging. Also, the GUI source code is available for custom modification.

The kit has a variety of features. There is an input and display section for the user to interact with the board. There is a remote communications section that has physical connections for Ethernet (RJ-45), RS-232 (DB9), and USB. There is a section with temperature sensors, a heating element, and a fan, enabling closed-loop thermal control. There is a power monitoring section with ATX connectors that can be easily monitored with the nice GUI. There is a prototyping section with a number of headers, including an Actel-standard 48-pin header and a Santa-Cruz standard header. There are two PCI card slots so a

customer can connect a standard PCI card for developing system management next to a payload application. Finally, there is on-board SRAM and Flash memory for extra memory capability during development. The FPGA logic and Flash can be programmed through the JTAG header. The off-chip memory can also be programmed through this header.

This development kit includes all hardware cables (including an optional FlashPro3), a software GUI, a demonstration design using the on-chip ARM7 processor, and all appropriate documentation, including user’s guides and schematics. It is useful for proprietary and standard applications, such as MicroTCA, ATCA, and IPMI.

Summary

System management can no longer be an afterthought to board design. MCU and FPGA complexities require that system management must be incorporated throughout every part of the design process. Today's current discrete implementations have a high design overhead associated with them, often requiring different components and board-level changes with each minor design iteration. A single-chip, configurable implementation is the ideal system management solution. Actel mixed-signal FPGAs offer an unprecedented level of integration, reducing component count, board space, and total system cost.

For more information, visit our website at **www.actel.com**

www.actel.com

Actel Corporation
2061 Stierlin Court
Mountain View, CA
94043-4655 USA
Phone 650.318.4200
Fax 650.318.4600

Actel Europe Ltd.
River Court, Meadows Business Park
Station Approach, Blackwatery
Camberley, Surrey GU17 9AB
United Kingdom
Phone +44 (0) 1276 609 300
Fax +44 (0) 1276 607 540

Actel Japan
EXOS Ebisu Bldg. 4F
1-24-14 Ebisu Shibuya-ku
Tokyo 150 Japan
Phone +81.03.3445.7671
Fax +81.03.3445.7668
www.jp.actel.com

Actel Hong Kong
Suite 2114, Two Pacific Place
88 Queensway, Admiralty
Hong Kong
Phone +852 2185 6460
Fax +852 2185 6488
www.actel.com.cn

55900069-0/12.06