Lowest Power FPGAs: IGLOO2 and SmartFusion2

July, 2013

Introduction to IGLOO2 FPGAs and SmartFusion2 SoC FPGAs

IGLOO2® FPGAs and SmartFusion®2 system-on-chip (SoC) FPGAs are differentiated from other FPGAs by their low power capabilities that enable orders of magnitude lower power operation for low duty cycle applications, refer to Figure 1 on page 3. Both device families include several important low power features:

- Industry's lowest static power
- Flash*Freeze real-time low power state
- ARM® Cortex™-M3 low power modes (SmartFusion2 SoC FPGAs only)
- SoC peripheral low power modes (SmartFusion2 SoC FPGAs only)
- Dedicated High-Performance Memory Subsystem to efficiently implement common bridging functions between 5 Gbps SERDES-based communications ports, external high-speed DDR memories, large on-chip SRAM or NVM memories and FPGA fabric. (IGLOO2 FPGAs and SmartFusion2 SoC FPGAs)

In systems that operate reactively or periodically, these devices can dramatically reduce power. Reactive operation is defined as being in a standby state waiting for some activity or event before initiating processing. The reactive system returns to the standby state after processing is complete. The following are examples of reactive systems:

- Patient health monitor alarm that is activated when a patient falls down
- Remote sensor initiating communication based on event detection

Periodic operation is defined as consisting of some activity that must be performed on a recurring basis. These systems perform some processing and then enter standby mode. At a fixed time interval, the system repeats the processing function and returns to standby mode. The following are examples of periodic systems:

- Many standard wireless protocols
- Patient heart rate monitor or similar that measures the pulse pressure periodically
- Remote sensor periodically measuring information

Saving energy in periodic and reactive systems is achieved by moving into a very low power state when processing is not necessary. For example, in SmartFusion2 FPGAs if the ARM Cortex-M3 is not required to operate during part of a data processing algorithm, perhaps because the FPGA fabric is implementing an autonomous function that doesn't require CPU intervention, the CPU can be put into a low power 'sleep' state until it is again needed for processing. In both device families the FLASH*Freeze mode can be used to dramatically reduce power in the FPGA fabric and the I/Os. The capabilities of Flash*Freeze mode, however, are completely unique in FPGAs.

Flash*Freeze technology enables the rapid stopping and starting of the FPGA fabric and related I/Os while preserving the state of the FPGA fabric and dramatically reducing power. The time to enter Flash*Freeze mode is approximately 100 µs; the time to exit Flash*Freeze mode is also approximately 100 µs. While in Flash*Freeze mode, the state of the FPGA is maintained so that upon exit, the device continues to operate where it left off. In the Flash*Freeze mode the FPGA fabric power consumption can be reduced to only 1 mW, orders of magnitude lower than the fabric power consumption of SRAM-based devices.

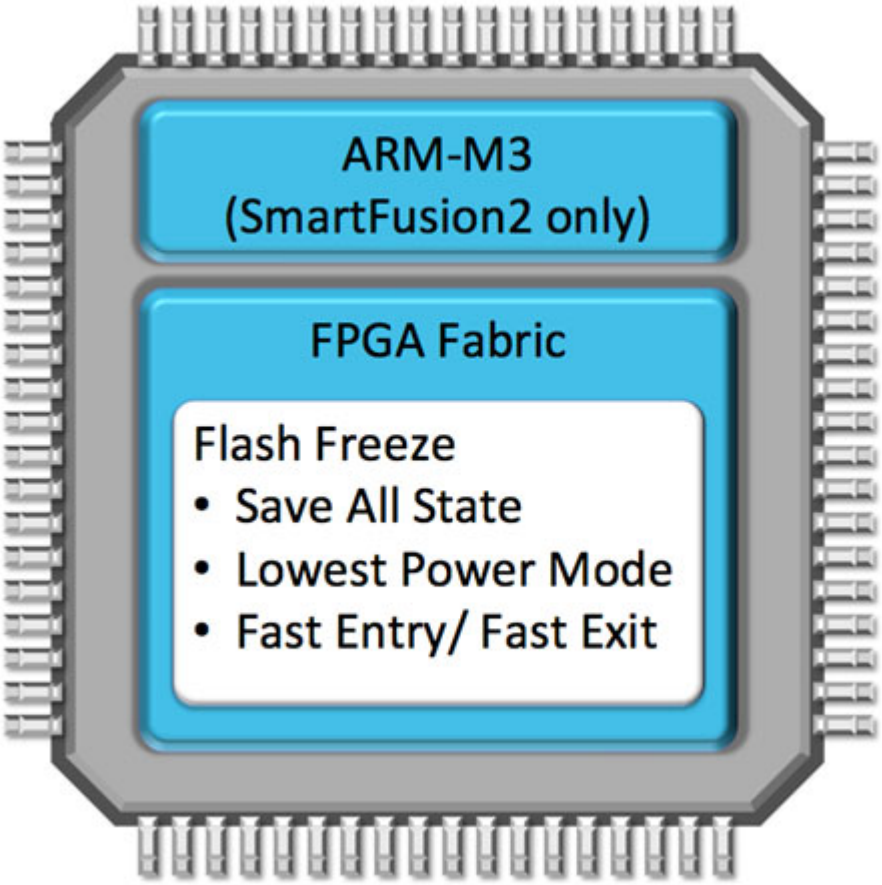

Figure 1: **SmartFusion2 and IGLOO2 Flash*Freeze-Features**

Flash*Freeze Operation

Figure 2 on page 4 shows an example of Flash*Freeze operation with a 10 ms duty cycle, as is used by commercial 3G Wideband Code Division Multiple Access (W-CDMA) based wireless communication. In W-CDMA systems, when the phone is in standby mode, some specialized circuitry wakes up every 10 ms and listens for a random access channel (RACH) signal, which indicates an incoming call. When no call is detected, the circuitry goes back to sleep, thus significantly saving battery life.

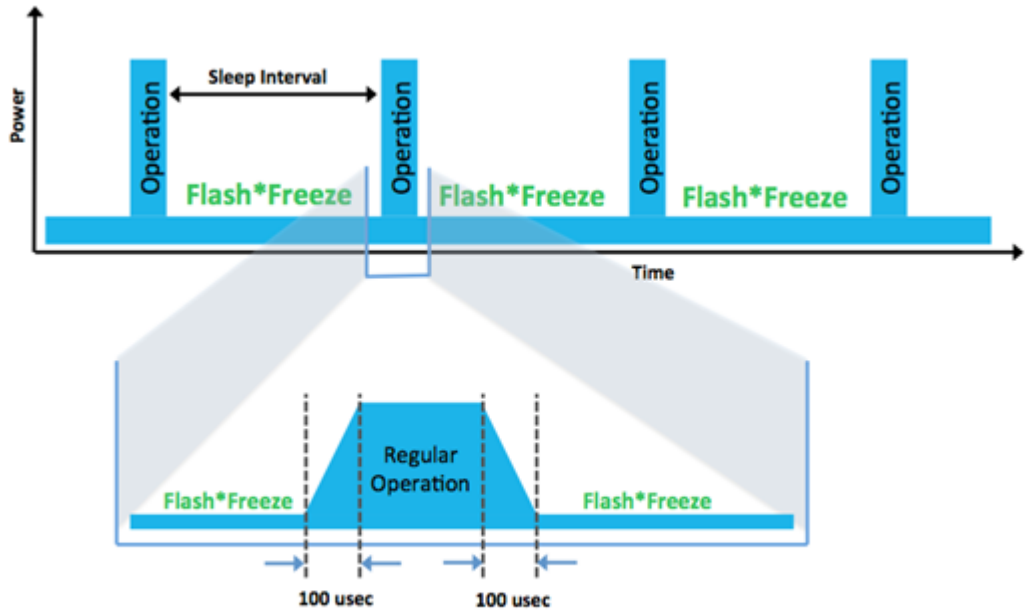

Figure 2: **System Moving from Flash* Freeze Mode to Regular Operation and Back**

In Figure 3, the total "on" time is 300 µs plus 100 µs to exit Flash*Freeze mode, and another 100 µs to enter Flash*Freeze mode. For the other 9.5 ms of the period, the FPGA is in Flash*Freeze mode, consuming very low power. During the “on” time of 300 µs, the SmartFusion2 device is capable of processing about 50 thousand M3 instructions. Both SmartFusion2 and IGLOO2 devices are capable of executing 50 thousand FPGA clock cycles and up to 10 million multiply-accumulate operations.[1] In this example, the resulting power reduction for the FPGA is an additional 20X over a device that does not have Flash*Freeze capability.

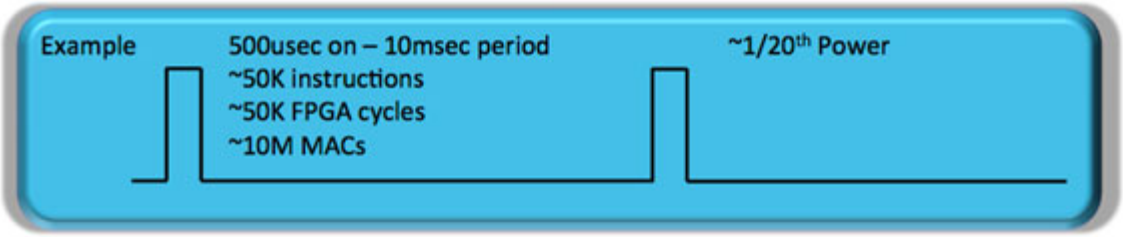

Figure 3: **10 ms Period Flash* Freeze Operation**

1. Clock frequency of 166 MHz is assumed for Cortex-M3 processor (in the SmartFusion2 implementation only), FPGA fabric, and Ethernet MACs.

Triggering Flash*Freeze Exit

Exiting Flash*Freeze mode can be initiated through I/Os (on both device families) or through the Cortex-M3 processor (on SmartFusion2 devices) using the special communications port. These methods send a message to the system controller to begin Flash*Freeze exit, refer to Figure 4.

Using I/Os to trigger Flash*Freeze exit can be done with Signature mode or Activity mode settings. In Signature mode, a selected set of I/Os is configured as inputs with predetermined 1 and 0 comparison states. Once entering Low power mode, every I/O designated as a signature I/O becomes input only. All other I/Os are tri-stated, held by bus hold, or weakly pulled-up/pulled-down. When the input signals for all the Signature mode I/Os match their predetermined 1 and 0 states, the system will initiate Flash*Freeze exit. All I/Os in Signature mode are compared as a single signature. In Activity Mode a selected set of I/Os are configured as inputs. When entering Activity mode, the value at the pin of the designated Activity I/O is latched. Any change on any of the Activity inputs will initiate Flash*Freeze exit. The device can also have a combination of one group of Signature mode inputs and a set of Activity mode inputs, either of which can initiate Flash*Freeze exit.

Signature and Activity mode triggers are particularly useful with IGLOO2 devices since the entire device is placed into the Flash*Freeze mode and an external signal is required to ‘wake’ the device. One simple example of using the signature mode to wake the device could be monitoring an address bus. Only if the specified address and command signals are observed would the device exit from Flash*Freeze.

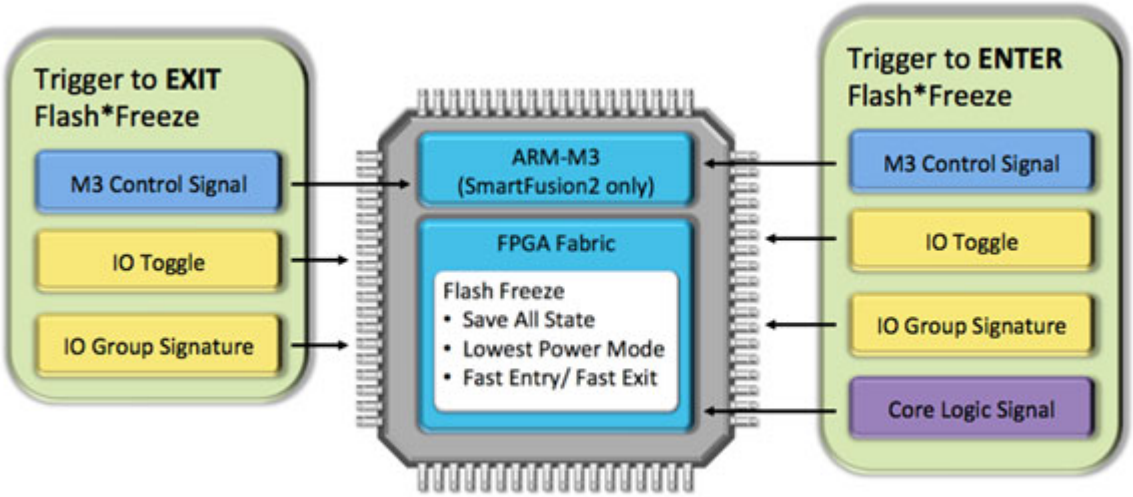

Figure 4: **Triggers to Enter and Exit Flash*Freeze Mode**

Initiating Flash*Freeze Entry

Flash*Freeze low power mode can be initiated from the Cortex-M3 processor, the I/Os, or a core logic signal in the FPGA fabric. This flexibility includes all the capabilities for exiting Flash*Freeze mode, but also adds the FPGA fabric as an additional method. Using the FPGA fabric, nearly any desired complex event or state can initiate Flash*Freeze low power mode.

Security with Flash*Freeze Mode

All methods of entering Flash*Freeze mode are controlled by the user during the design. In secure systems, all methods of entering Flash*Freeze mode should be controlled from within the FPGA and should not utilize any external signaling. This keeps external attackers from utilizing Flash*Freeze mode to disable the system.

Flash*Freeze mode can also be used within security systems as a punishment where several pass-key checks have failed—similar to smartphone time-out when too many password entries have failed. This is one method to dramatically decrease the number of pass-key attacks that can be performed on a system, thus mitigating brute force attacks.

Flash*Freeze Applications

Communications

In low power communication systems, using periodic bursts of communication reduces power. This eliminates constant power in the amplifiers and the rest of the system.

Sensor Networks

Sensor networks include both low power communications and active sensors to perform distributed measurements. The active sensors can be turned on periodically (traffic image, weather sensor) or can turn on in response to an event (earthquake). After the measurement has been taken, the information is uploaded as an information burst and then the equipment goes back to sleep using Flash*Freeze mode.

Medical Equipment

Many types of medical equipment are used for monitoring patient health. These types of systems have relatively low sampling rates and can therefore utilize periodic operation as a means to minimize power consumption. This is especially useful in portable medical equipment.

Lowest Static Power FPGAs

By using on-chip non-volatile Flash memory for configuration, IGLOO2 FPGAs and SmartFusion2 FPGAs don’t require an external configuration device. SRAM-based FPGAs require an external configuration device, which increases system power requirements. Additionally, during FPGA start-up, SRAM FPGAs create a large power ‘spike’ and this contributes to an additional power drain. IGLOO2 FPGAs and SmartFusion2 SoC FPGAs, with on-chip Flash configuration, have significantly lower static power. For example, a 50K gate device has only 10 mW of static power, a fraction of the leakage current of a comparable SRAM-based FPGA.

Industries Lowest Power FPGAs

SmartFusion2 SoC FPGAs and IGLOO2 FPGAs are the lowest power devices in the industry. They share the inherently low static power of on-chip Flash configuration and the innovative Flash*Freeze mode that can reduce FPGA fabric power consumption to only 1 mW and takes no more than about 100 µs to enter or exit. Additionally, SmartFusion2 SoC FPGAs implement power saving modes in the ARM Cortex-M3 processor and SoC peripherals. The Cortex-M3 core has Sleep and Deep Sleep modes that allow the device to be put into a low power state. Sleep mode stops the processor clock to reduce power consumption, and Deep Sleep mode also stops the processor PLLs and flash memory. The SoC peripherals can be shut down individually if they are not used in a design to reduce power consumption further.

SmartFusion2 SoC FPGAs are the lowest power SoC FPGAs with an integrated ARM Cortex-M3 processor and are ideal for use in any power sensitive application in the industrial, medical, military, avionics and communication markets. The IGLOO2 on-chip dedicated High-Performance Memory Subsystem (HPMS) features large SRAM and NVM blocks, dual high-performance DMA controllers, dual external DDR Memory controllers, dual Fabric Interface Controllers (to easily connect to fabric based functions) and the intelligent AHB Bus Matrix used to efficiently interconnect each HPMS block. This makes IGLOO2 devices the lowest power FPGAs for data transfer and data storage bridging oriented designs. IGLOO2 devices are thus ideal for use in these types of power sensitive applications in the industrial, medical, military, avionics, and communication markets.

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo CA 92656 USA
Within the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

Microsemi Corporation (NASDAQ: MSCC) offers a comprehensive portfolio of semiconductor solutions for: aerospace, defense and security; enterprise and communications; and industrial and alternative energy markets. Products include high-performance, high-reliability analog and RF devices, mixed signal and RF integrated circuits, customizable SoCs, FPGAs, and complete subsystems. Microsemi is headquartered in Aliso Viejo, Calif. Learn more at **www.microsemi.com**.

55900157-0/7.13