White Paper

PolarFire™ Non-Volatile FPGA Family Delivers Ground Breaking Value: Cost Optimized, Lowest Power, SEU Immunity, and High-Security

10 2017

Microsemi
Power Matters.™

Background: Market Dynamics

The central nervous system serves as a framework for characterizing the rapidly evolving global market for electronic hardware. The cloud is the brain, connected devices are receptors, and the transport network is the nerves that connect the devices to the brain. The cloud is realized by data centers, the transport network by fiber and microwave backhaul, and connected devices by a diverse array of electronic equipment, including sensors, actuators, and mobile devices.

There are two megatrends that are having a deep impact on how this market evolves, and so are driving innovation.

- The end of Moore’s law and the related Dennard scaling.
- The exponential growth rate of connected devices. Mobile phones kicked off the first wave of this megatrend and the Internet of Things (IoT) represents the second wave. A new element of the current wave is that these devices, which include cars, are increasingly playing a role in providing safety for the humans they interact with.

Both megatrends have elevated power to a zeroth concern of system designers in all domains. In the data center, the constraint is the fixed amount of power supplied to the data center. In the past, the data center operator could count on Moore’s law and Dennard scaling to provide a doubling of compute capacity of the fixed investment by simply upgrading to the latest generation of processors. At the device level, system power is limited either by batteries (for example, cell phones) or a limited supply such as power over Ethernet. In between, the transport layer and the gateways that feed it are often deployed in fanless enclosures in outdoor environments. In these form factors, available power is not the problem, removal of heat is the problem. The exponential increase in connected devices and the data they produce is driving up the capacity requirements (required power) for both the transport layer and the data centers.

Several solution trends have emerged to address the needs and constraints of this market dynamic. Compute workloads are being pushed from the data center out to the gateways and devices, lowering the power requirements of the data center and transport network while also reducing latency in the system. Heterogeneous compute paradigms (the addition of accelerators and GPUs to the traditional microprocessor based computer platforms) are being deployed in data centers, gateways, and devices to compensate for the loss of gains in processing capability that arose because of Moore’s Law and Dennard scaling. Gateways and devices, which are exposed to physical attacks, are deploying hardware-level security. Finally, the increased exposure of humans to connected devices increasingly requires system designers to consider safety and reliability requirements.

Figure 1 Market Dynamics

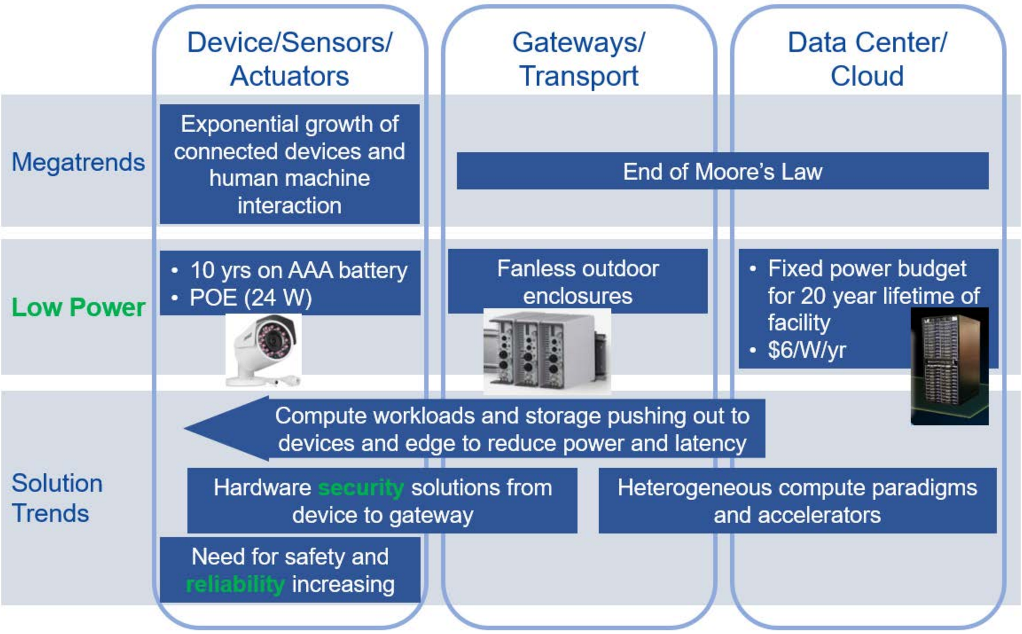
Device/Sensors/ Actuators
Gateways/ Transport
Data Center/ Cloud
Megatrends
Exponential growth of connected devices and human machine interaction
End of Moore's Law
Low Power
• 10 yrs on AAA battery
• POE (24 W)
Fanless outdoor enclosures
• Fixed power budget for 20 year lifetime of facility
• $6/W/yr
Solution Trends
Compute workloads and storage pushing out to devices and edge to reduce power and latency
Hardware security solutions from device to gateway
Heterogeneous compute paradigms and accelerators
Need for safety and reliability increasing

This document describes how the new PolarFire FPGA family from Microsemi can help system designers deploy these solution trends, particularly in the transport layer, associated gateways, and connected devices. Built on 28nm Silicon-Oxide-Nitride-Oxide-Silicon (SONOS) technology, the FPGAs provide market-leading reliability, low-static power, and hardware security. The FPGA fabric includes several innovations that contribute to the FPGA's leading low-power capabilities. The FPGA's low-power transceivers operate between 1 Gbps and 12.7 Gbps, and are ideally suited for the gateway applications' bandwidth requirements. Finally, significant enhancements to Microsemi's market-leading security capabilities have been deployed on the FPGA, including an advanced crypto-processor.

PolarFire SONOS Technology

The PolarFire FPGA family uses SONOS non-volatile (NV) technology on a 28nm technology node. This provides a cost advantage over SRAM-based FPGAs at the same or even smaller node and relative to Microsemi's previous generation FPGAs using floating gate NV technology (65nm and older). The use of the 28nm technology node scaling factor enables a significant cost and performance advantage relative to Microsemi's previous technology nodes. The cost advantage comes from the scaling factor going from 65nm to 28nm of ~50%. The basic transistor performance advantage is ~2.5x, between 65nm and 28nm, using an inverter propagation delay for comparison.

How SONOS and Floating Gate NV Technologies Work

The device is programmed when an electronic charge is transferred to the charge storage layer to create a bias affecting the NV transistor's characteristics. In the case of a negative charge, it acts as if the metal oxide semiconductor (MOS) transistor gate has a negative bias applied, so the device will be strongly "off." In the case of a positive charge (depletion of electrons, or additional "holes"), the gate is positively biased, so the device will be strongly "on."

The Floating Gate Non-Volatile Device Used in Previous Flash Architectures

The floating gate technology requires 17.5 V with large charge pumps that consume a substantial die area. The floating gate technology uses a double-poly transistor stack with a conductive Poly-1 layer as the charge storage element, called the floating gate. The bottom oxide thickness is critical for both preventing charge loss due to defects and determining the programming voltage. A thicker bottom oxide prevents charge loss due to oxide defects but requires a higher programming voltage. Microsemi's FPGA products use a relatively thick bottom oxide for high reliability, preventing charge loss due to oxide defects, and 17.5 V to program.

The following illustration shows a detailed description of the floating gate non-volatile device.

Figure 2 Floating Gate Non-Volatile Transistor

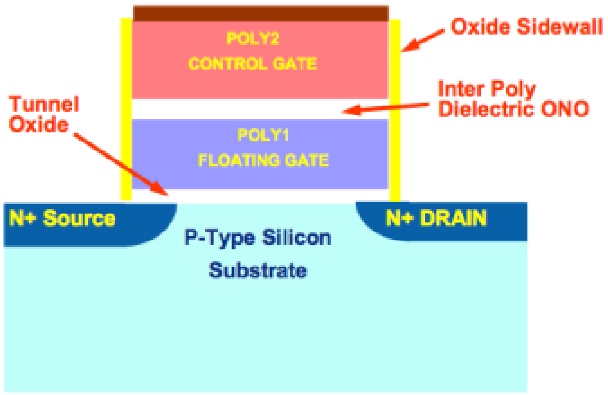

PolarFire SONOS Non-Volatile Device

The SONOS technology requires only 7.5 V for programming, so charge pumps can be smaller. This technology enables a smaller die size and contributes to a more cost-effective device.

The SONOS device uses a single poly transistor stack with a non-conductive Nitride dielectric layer (silicon-nitride, Si3N4) as the charge storage element. The advantage is that if a defect exists in the bottom oxide, only a very small amount of charge will be lost in proximity to the defect. Most of the stored charge stays intact where it is because the stored charge is non-mobile in the insulating Nitride layer. This allows the use of a thinner bottom oxide that can be programmed with lower programming voltages (~7.5 V) and smaller charge pumps, compared to the floating gate technology.

The following illustration shows the SONOS transistor.

Figure 3 SONOS Technology

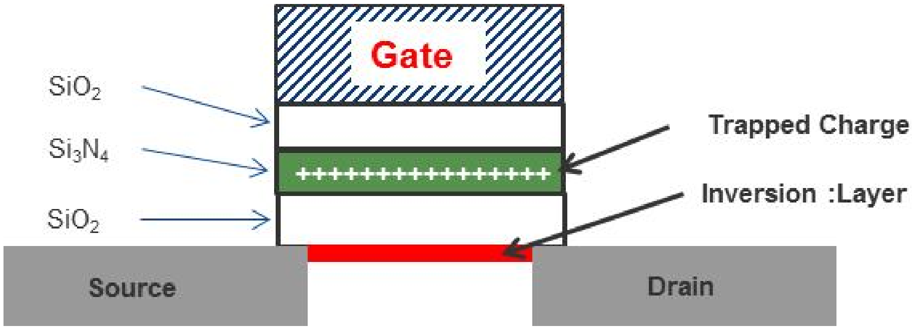

Though SONOS requires some extra process steps compared to an unenhanced CMOS manufacturing process, it uses fewer transistors than an SRAM memory element, so is very cost-competitive.

Reliability

The 28nm SONOS NV technology uses a push-pull cell containing an N-channel and a P-channel NV device. The NV devices are not in the data-speed path and are only used to control a standard transistor used as the data-path switch. This provides a large functional advantage because any variation in the NV device threshold voltage (V_t) does not change the switch conductance.

A simple description of a push-pull cell is that the N-ch and P-ch devices are stacked in a series with each other between the power and ground rails, with one in the "on state" and one in the "off state." The N-ch and P-ch NV devices compete against each other to control the gate of the switch transistor. The On device will overpower the Off device and drive the gate of the switch device to a high or low voltage (depending on which NV transistor is on and off), thus putting the switch in either an on or off state. If either of the NV devices is a weak bit (that is, at the lowest Vt limit allowed), the other bit still holds the correct state. This acts as a built-in quasi redundancy, because one NV device can be weak and no performance degradation will occur over the life of the product.

In the example of the operation for the state 1 case (as seen in Figure 1), the P-ch device is on and the N-ch device is off. The P-ch device passes its source voltage to the switch (gate voltage >= the Vt of the switch transistor), and so the switch device is on (the channel is highly conductive).

All PolarFire FPGA family dies are screened for full NV retention lifetime before shipment. The screening is done as part of the standard wafer-sort manufacturing flow.

Low Power

A significant reason why the PolarFire FPGAs consume such low power compared to competing FPGA technologies comes from the SONOS NV FPGA configuration cell. The following illustration shows two schematics of the NV cell, highlighting the different programmable "configuration" states that control the FPGA data signal path. There is an off state where the datapath is switched off and an on state where the datapath is switched on.

Figure 4 Low Power

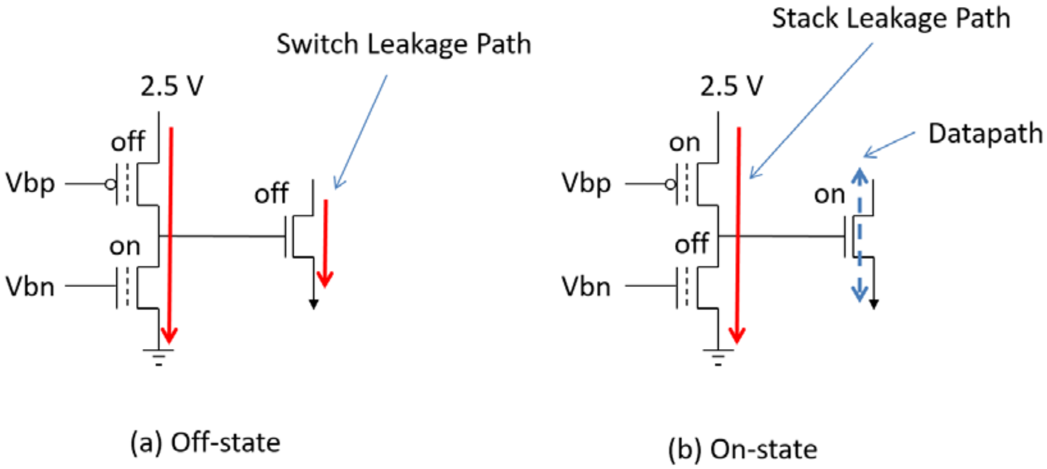

Consider the stack and switch leakage paths. In the "stack" leakage path, one of the two NV elements is always programmed into a very deep off-state condition. Using "on state" as an example, the N-ch NV element is in the off state with its Vt shifted ~0.5 V above a normal transistor Vt, consequently the leakage will go down to a negligible level. The leakage of the NV stack is extremely low—much lower than the leakage of a standard CMOS transistor stack. In addition, there are fewer transistors in an NV configuration memory cell than in an SRAM memory cell.

The "switch" leakage path is the leakage across an "off-state" switch—the "FPGA logic signal path" leakage. The switch device is a high-voltage device and has been optimized to provide much lower leakage than a standard transistor.

Flash*Freeze™ Mode

Another low-power advantage of using our NV technology is a power-saving feature called Flash*Freeze mode. In this mode, the product can be put into a state that turns the supply voltage off to the configuration memories in the FPGA logic block while saving the user's state in low-power latches, thus lowering the standby power by approximately two-thirds. This is a unique feature enabled by the usage of the NV configuration cell that is not possible with volatile FPGA technologies. The NV cell will retain its state after power has been turned off to the device, allowing the FPGA to return to normal operation without reconfiguration.

The Live-at-Power Up ("Instant On") and "Single Chip" Features

An advantage of the NV technology is that there is no need to reload the FPGA design code when power is returned because the FPGA logic configuration cell retains its state after power down. Thus, there is no need for an external boot PROM and the programmed FPGA is fully functional as a single chip. Also, the boot time is very fast because there is no large transfer or decryption of data that must occur before the FPGA is usable. The millions of configuration cells are directly controlling the corresponding switch transistors as soon as power is applied.

Single-Event Upset (SEU) Immunity

The FPGA logic configuration is SEU immune because of the non-volatile technology, unlike the configuration memory in SRAM-based FPGAs, which can flip state due to neutron hits.

Configuration memory upsets are especially problematic because the configuration memory must remain static and error free during all the operating hours of the device for correct FPGA operation. Any upset will be persistent until the device is powered-down or the cell is reprogrammed correctly. If an upset occurs in the erroneous state, the logic or routing of the FPGA fabric will be wrong, potentially causing not just a single wrong data value, but a string of wrong results until it is fixed. This may require a full system reboot.

In the PolarFire FPGA family, the SONOS NV charge is stored in the nitride dielectric, which is not susceptible to charge loss from neutron hits. The net effect is that Microsemi FPGAs are magnitudes more reliable than competing SRAM FPGAs.

PolarFire FPGA Fabric

In designing the FPGA programmable logic fabric, we aimed to meet mainstream performance requirements with minimal power and cost. Microsemi's PolarFire FPGAs typically consume one-tenth the static power of competing SRAM FPGAs, and half the total power. Some attributes of the PolarFire's FPGAs (such as the non-volatile configuration memory) directly reduce power, while power reduction in an indirect effect of reducing die area in other cases.

Selection of LUT-4 for the Logic Element

6-input LUTs can provide some speed benefits, but 4-input LUTs are the better choice for a power- and cost-optimized FPGA like Microsemi's in a modern process technology. It has been well-established that 4-input LUTs can make more efficient use of a die area than 6-input LUTs. A given user design can be implemented with less silicon area using a 4-LUT architecture than using a 6-LUT architecture. One contributing factor is that a 6-input LUT requires 4x more configuration memory bits (64 versus 16) but can accommodate only about 1.6x as much logic as a 4-input LUT. This traditional observation applies even more strongly to advanced fabrication technologies because SRAM configuration memory has not scaled as fast as ordinary logic, due to the need to mitigate the risk of SEUs. The PolarFire SONOS configuration cell is immune to SEUs.

Consider a PolarFire FPGA cluster of twelve 4-input LUTs versus a cluster of eight 6-input LUTs. The total logic capability of the cluster (that is, the amount of user logic that the cluster can accommodate) is similar in each case. The larger fan-in of the 6-input LUT means fewer levels of logic may be traversed by the critical path within each cluster, potentially reducing the total contribution of intra-cluster delay to the critical path. However, from the outside, the two clusters appear similar; they have a similar typical number of incoming and outgoing signals, and so the total length and delay contributed by the inter-cluster wiring is similar in both cases.

The following illustrations show clusters with different LUTs, but similar numbers of incoming and outgoing signals.

Figure 5 Cluster of Twelve 4-Input LUTs

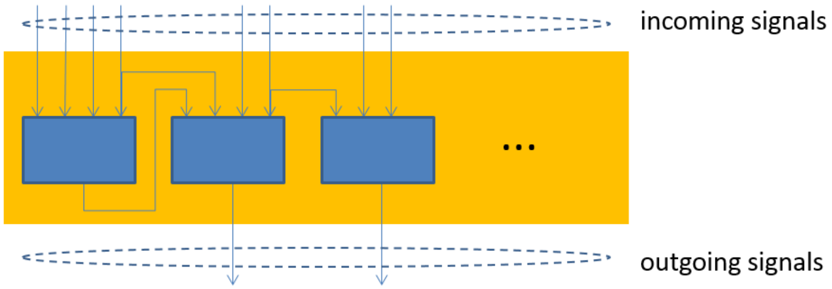

Figure 6 Cluster of Eight 6-Input LUTs

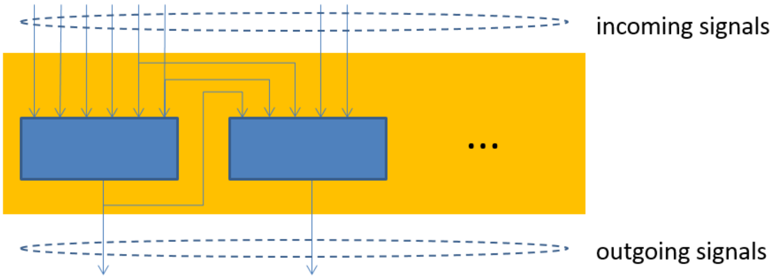

As process technology has progressed from 65nm to 28nm and beyond, the delay of wiring has come to dominate logic delay, due to poor scaling of metal wire and via resistance. To some extent, this can be mitigated by widening the wires, but that adds to the die area and cost. So, with each succeeding generation of process technology, inter-cluster delay becomes a significant contributor to the critical path, and the speed advantage of 6-input LUTs diminishes.

The PolarFire FPGA family provides rapid direct connections between nearby LUTs. This can reduce intra-cluster delay, especially in conjunction with advanced synthesis and placement algorithms. Certain logic functions (such as MUX trees) greatly benefit from the direct connections.

Clock Dynamic Power

Clocks can be a significant contributor to dynamic power in applications targeted by PolarFire FPGAs, often consuming nearly as much power as the rest of the routing and logic. For this reason, clocks in PolarFire FPGAs were designed to conserve power. We allocated more area to clock wires to space them further apart, significantly reducing their capacitance and dynamic power. Flip-flops were designed to minimize clock power. Clock gating is provided at two levels in the clock tree, as well as for each individual flip-flop, to avoid wasting power on unused branches. As a result, clock power in PolarFire FPGAs is less than half that of competing 28nm FPGAs, averaged over a suite of designs (as shown in the following illustration).

Figure 7 Clock Power Comparison on Various Benchmarks

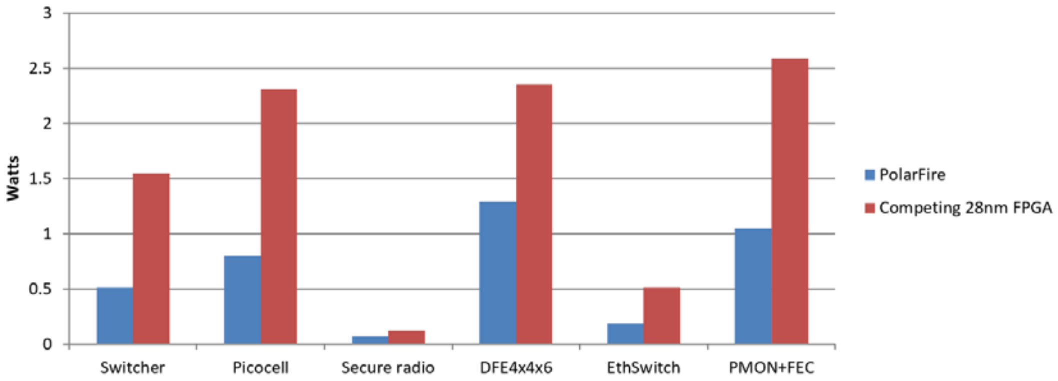

Choice of Operating Voltage

The PolarFire FPGA family's power-performance tradeoff has been carefully optimized for a 1.0 V core logic supply, somewhat less than the 1.05 V nominal voltage for the UMC 28nm process on which it is manufactured. Customers desiring extra speed still have the option to use the full 1.05 V supply.

Math Block

The PolarFire FPGA provides a math block supporting 18-bit multiply-accumulate operations (as seen in Figure 8). The following are improvements from the previous-generation IGLOO™2 FPGA family.

- Provision of a pre-adder with a full 19-bit result. This eliminates the need for fabric adders when implementing symmetric FIR filters, saving power.
- Provision of an input value cascade chain. This reduces the need for fabric registers when implementing systolic FIR filters, again saving power.
- Accumulator widened to 48 bits.

In addition to the 18-bit × 18-bit multiplication mode, the math block supports reduced precision 9-bit operations. The PolarFire math block supports two independent 9 × 9 multiplies with no requirement for a common factor. Unlike another competing math blocks, which can exchange two 18 × 18 multipliers for three 9 × 9 multipliers, PolarFire FPGA can exchange one 18 × 18 multiplier for two 9 × 9 multipliers, a 33% improvement.

The PolarFire math block also supports a unique 9 × 9 dot-product mode (as seen in Figure 9). This is ideal for use in image processing and convolutional neural networks (CNNs). Compared to independent 9 × 9 multipliers, the dot-product operation reduces power in the following ways.

- No need for a separate fabric adder to sum the two products.
- The pre-adder is fully supported, allowing efficient implementation of symmetric 9-bit FIR filters or 2-D convolution.
- All four factors are independent—designers of CNNs don’t rely on complex weight-sharing or input-sharing schemes to maximize resource and power efficiency.

The following illustration shows a simplified diagram of a math block dot-product mode.

Figure 8 Math Block Dot-Product Mode

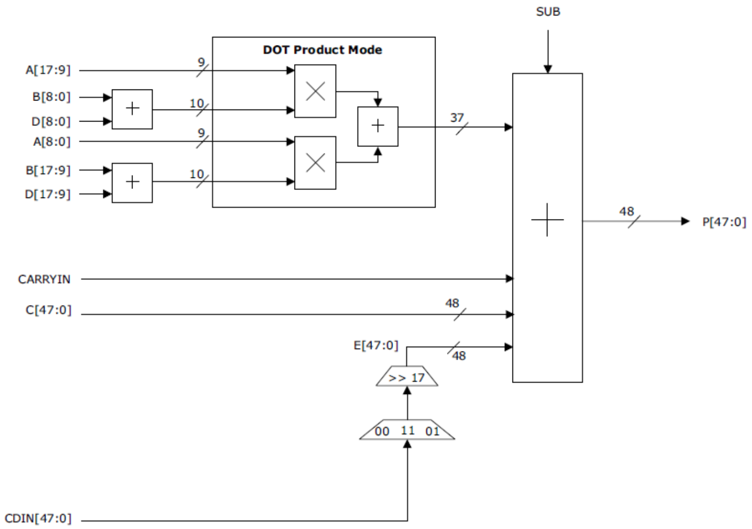

PolarFire FPGA Transceivers

PolarFire FPGAs include up to 24 high-speed full-duplex transceiver channels with unique features unavailable in competing devices. The SerDes transceiver was specifically optimized for PolarFire mid-range FPGAs, supporting baud rates from 250 Mbps to 12.7 Gbps, covering the full SDI range of applications. One primary advantage of optimizing for this range over downgrading a higher-speed Serd-Des adapted from a high-end FPGA, is that it has significantly lower power at all baud rates compared to transceivers on competitive mid-range FPGAs.

Exceptionally Low Transceiver Power

The transceivers used in PolarFire FPGAs have the lowest power consumption of any comparable-performance transceiver used in any FPGA. The following illustration shows a comparison of the estimated power consumption for several FPGA transceivers at different baud rates.

Figure 9 Power Comparison of PolarFire FPGA Transceiver to Competing SRAM FPGAs

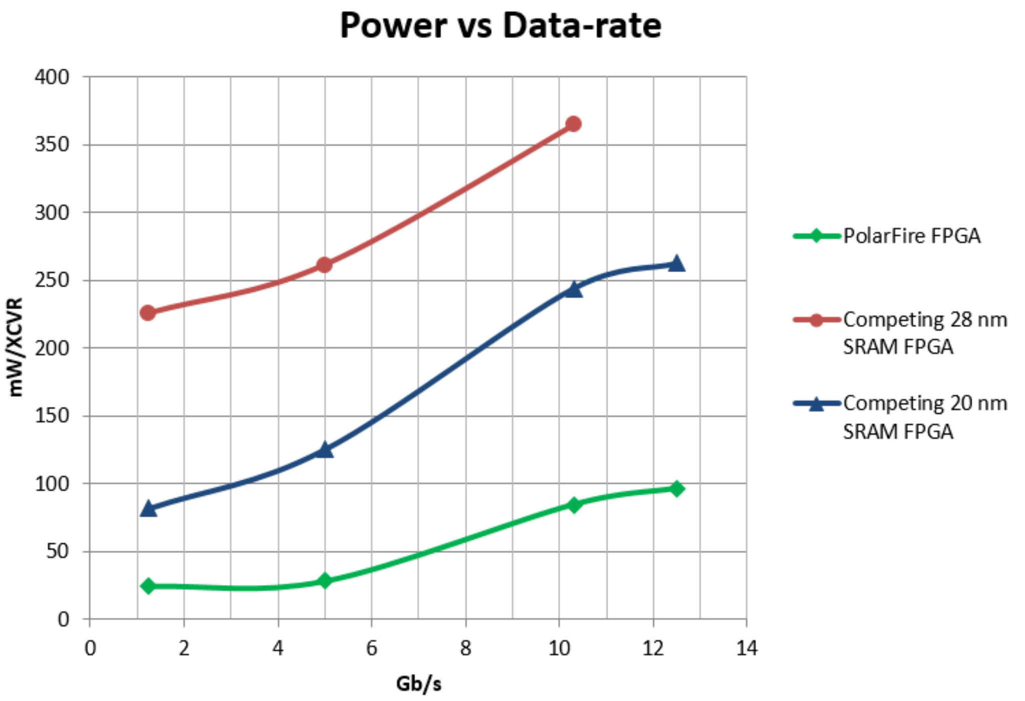

Transceivers in competitive FPGAs use more than twice the power of PolarFire transceivers and from 3x to 10x more estimated power at most baud rates. Low power consumption is achieved in the PolarFire FPGA transceivers through a unique combination of architectural decisions.

- The transceiver is implemented using a half-rate architecture. High-speed logic is clocked on both edges to keep the high-speed clock networks at half frequency.
- Low transmitter jitter is achieved without the use of high-power logic styles such as CML.
- The transmitter uses an LVDS-style output stage that has good noise immunity and uses less power than other driver types.

- PolarFire FPGAs have the unique ability to shut down the decision feedback equalizer (DFE) and eye monitor when not in use. This is a major power advantage, especially at lower data rates.
- The transmit PLLs use a highly-shared architecture for significant power savings; fewer transmit PLLs used in the device results in lower power.

Transceiver Features

PolarFire FPGAs have 1–6 quad-channel transceivers, for a total of up to 24 SerDes channels, depending upon the family member and device package. To see the exact number of transceivers supported in each device/package combination offered, see the device datasheet. Two second-generation PCIe endpoints/root ports are implemented in hardware, supporting a memory mapped AXI4 with built-in DMA (available in every PolarFire FPGA).

The following illustration shows the SerDes features.

Figure 10 Transceiver Architecture and Features

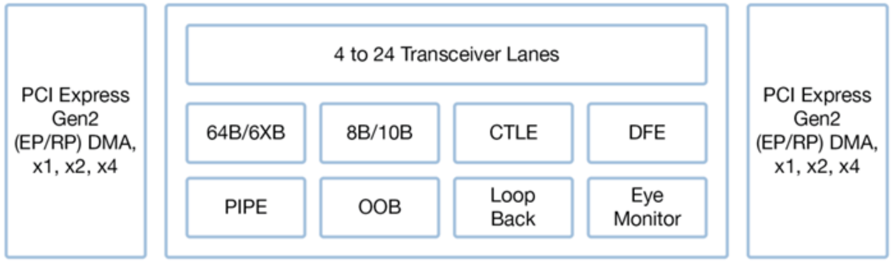
PCI Express Gen2 (EP/RP) DMA, x1, x2, x4
4 to 24 Transceiver Lanes
64B/6XB
8B/10B
CTLE
DFE
PIPE
OOB
Loop Back
Eye Monitor
PCI Express Gen2 (EP/RP) DMA, x1, x2, x4

Transceiver Special Features

The PolarFire FPGA transceivers have many special features that enable use with industry-standard protocols in mid-range FPGA applications.

Equalization Features

The following equalization features are supported by the PolarFire FPGA.

- The transmit channel has both pre- and post-tap feed-forward equalization (FFE) de-emphasis.
- The receive channel utilizes continuous time linear equalization (CTLE).
- The receive channel has a 5-tap decision-feedback equalizer (DFE).

Together, these features allow longer distances and/or the use of low-cost materials in printed circuit boards and backplanes.

Phase-Locked Loop (PLL) Features

The capabilities and features of the low-power PLL technology used for transmitting PLL and general purpose user-PLLs provides many user benefits.

- Fractional-N architecture, providing flexibility in clock and baud-rate selection.
- Spread-spectrum capability to easily meet maximum radiated emission requirements.
- Removal of transmit jitter up to the top baud rate of 12.7 Gbps, enabling higher bandwidth than competing solutions, and support for Sync-E (1 GbE to 10 GbE).
- Support for unique and independent transmit and receive rates in each transceiver channel.
- Unique CDR PLL capability to support burst mode protocols (such as GPON and XGPON) without the use of oversampling (lower power).

Differentiating Debug and Test

The PolarFire FPGA transceiver has all the expected debug and test features including built-in pseudo-random binary sequence (PRBS) generators and checkers and IEEE 1149.6 “AC JTAG” support for non-DC-coupled signals. Another standard feature is a built-in eye monitor to allow optimal in-situ tuning of FFE and CTLE parameters and problem diagnosis.

The PolarFire FPGA eye monitor can be used in-application to identify the eye margin of the receiver. The following illustration shows the 10.3125 Gbps receive eye measured by on-chip eye monitor with a 26” backplane (FR-4 strip-line) and 7’ line-card (FR-4 micro-strip).

Figure 11 Receiver Eye Diagram

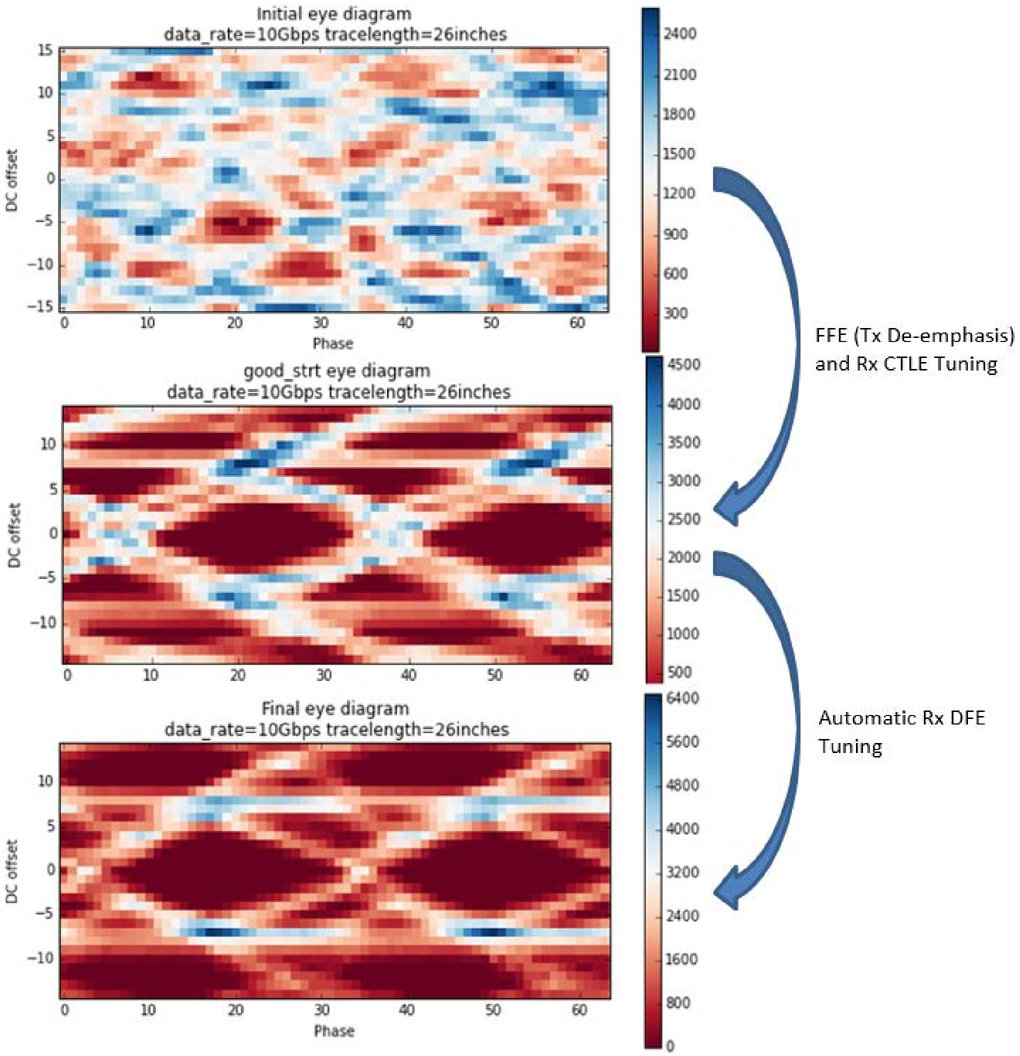

Measuring the receiver eye diagram improved using FFE, CTLE, and DFE.

Transceiver Summary

The PolarFire FPGA transceivers have the best-in-class implementation for any mid-range FPGA. They have unequaled performance and features at a power consumption point several times lower than the transceivers in competing FPGAs.

Best-In-Class Security

As FPGA capacities become larger, the value of the design IP has become greater. Additionally, the threat landscape has become decidedly more hostile with new vulnerabilities (such as differential power analysis being discovered) and a greater number of more aggressive, knowledgeable, and well-funded adversaries. Cost pressures are driving manufacturing into less trusted locations. These factors are affecting the security of the devices and their supply chain, the user's design IP and the FPGA configuration process, and any sensitive information the end system may be called upon to process.

The following illustration shows some of the potential threats in your supply chain.

Figure 12 Potential Threats to Systems Utilizing FPGAs

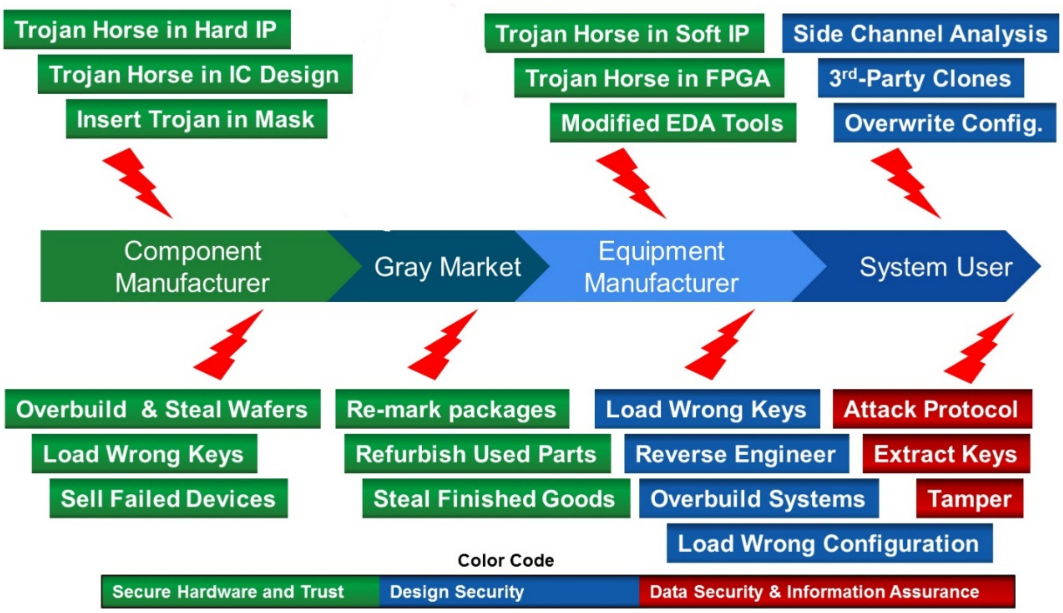

The techniques and features required to provide the level of security required in today's high-stakes, high-threat environment include a hardware roots of trust, strong cryptography coupled with top-notch key management at every stage, and devices with built-in passive and active countermeasures to protect against tampering. PolarFire FPGAs provide comprehensive best-in-class security at all stages of their life cycle.

Trusted Hardware

The foundation for all information security is trusted hardware. If you can't trust the hardware to do what it is supposed to do (that is, nothing malicious), then the security war is lost at the first battle. Microsemi has taken extraordinary steps to secure the supply chain so users of the PolarFire FPGAs can be assured that they are working with trustworthy devices.

Microsemi provisions the PolarFire FPGA family devices with cryptographic keys and certificates that can be used to verify each FPGA is authentic. The provisioning process is done securely, using FIPS140-2 level 3 certified hardware security modules (HSMs) placed at its wafer probe and package test facilities, preventing the possibility of rogue insiders subverting the process. All secret keys are encrypted and authenticated during transit and encrypted while stored, only being generated and used within the secure hardware boundary of an HSM or the FPGA. All security protocols (including key verification) are designed to prevent monitoring, man-in-the-middle, and replay attacks from succeeding. For example, Microsemi's PolarFire FPGAs, unlike the competitors', have a built-in cryptographic-grade true random number generator (TRNG) that is used to ensure each protocol session is "fresh" and not a replay of a previous session.

Besides public data, such as a unique-per-device serial number, only Microsemi FPGAs are provisioned with secret symmetric and asymmetric keys and an X.509-compliant public key certificate. These can be used to securely identify the device and to load confidential user keys and the initial user security policy. The X.509 certificate is issued and digitally signed by Microsemi, only for verified devices that pass the test, which prevents counterfeit devices or devices that failed and were supposedly scrapped (also known as floor sweepings) from entering the supply chain undetected.

The following illustration shows the secure provisioning of each PolarFire FPGA with a unique serial number, keys, and its X.509 public key certificate.

Figure 13 Device Certificate Chain of Trust

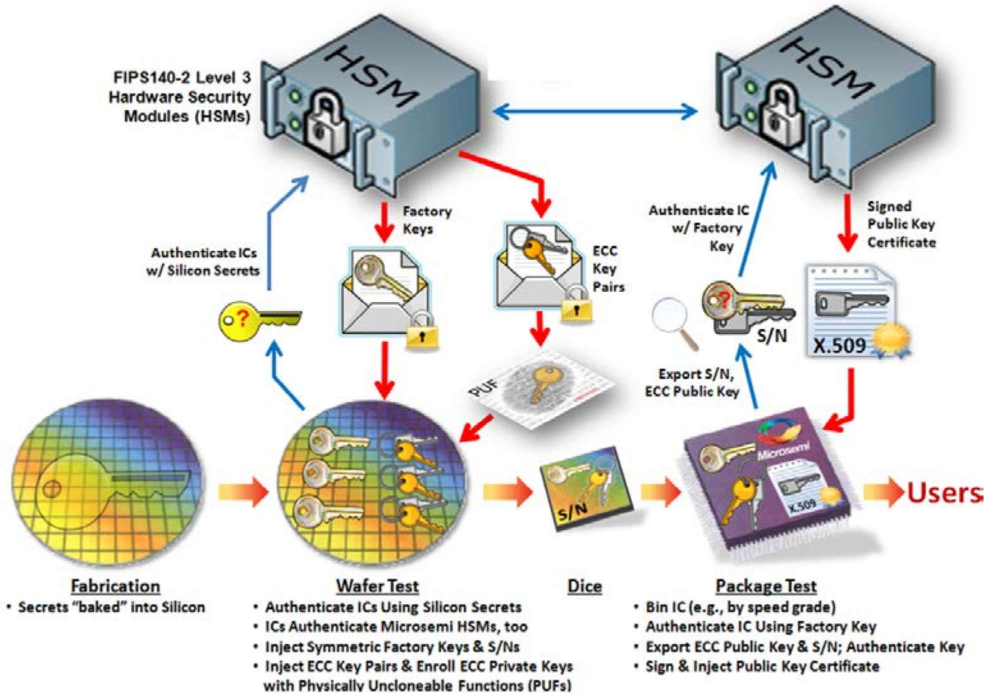
FIPS140-2 Level 3
Hardware Security
Modules (HSMs)
HSM
HSM
Factory Keys
ECC Key Pairs
Authenticate ICs w/ Silicon Secrets
Authenticate IC w/ Factory Key
Signed Public Key Certificate
PUF
S/N
X.509
Export S/N, ECC Public Key
S/N
Microsemi
X.509
Users
Fabrication
• Secrets "baked" into Silicon
Wafer Test
• Authenticate ICs Using Silicon Secrets
• ICs Authenticate Microsemi HSMs, too
• Inject Symmetric Factory Keys & S/Ns
• Inject ECC Key Pairs & Enroll ECC Private Keys with Physically Uncloneable Functions (PUFs)
Dice
Package Test
• Bin IC (e.g., by speed grade)
• Authenticate IC Using Factory Key
• Export ECC Public Key & S/N; Authenticate Key
• Sign & Inject Public Key Certificate

Design Security

Design security is the protection of the user's IP, which is used to configure the PolarFire FPGA. This not only includes keeping the IP confidential during transit but also ensuring that it is authentic, coming unmodified from an authorized source, and that it is used only in the ways the user intended (for example, to program only a fixed number of devices and no more).

It is certain that manufacturing and field locations provide greater exposure to malicious activity than the user's design center did. PolarFire FPGAs utilize the keys and certificate provisioned by Microsemi, plus other techniques—such as patented countermeasures to differential power analysis

(DPA)—to provide best-in-class FPGA design security in both the manufacturing and field environments.

Side Channel Analysis (SCA)

SCA is a class of techniques used to extract secrets, such as cryptographic keys, from an electronic system by monitoring information leaked through unintentional side channels. Many SCA techniques were discovered by Paul Kocher and his associates at Cryptography Research, Inc. (CRI, now a division of Rambus) in the late 1990s. The following illustration shows the registered trademark of CRI in the United States and other countries, used under license.

Figure 14 Licensed DPA Security Logo

Today, over seven billion integrated circuits are produced per year under license to Rambus/CRI's DPA patent portfolio. Based on an independent assessment, CRI granted Microsemi use of the "DPA padlock" security logo in conjunction with the design security protocols of SmartFusion™2 and IGLOO2 FPGA families. No other FPGAs on the market have been similarly certified. It is anticipated that the PolarFire FPGA family will easily pass this certification also, using Microsemi's second-generation DPA-resistant design security technology.

Physically Unclonable Function (PUF)

Microsemi introduced the Quiddikey®-Flex SRAM-PUF technology (licensed from Intrinsic ID, BV) in the SmartFusion2 and IGLOO2 FPGA families. PolarFire FPGAs use the latest incarnation of Quiddikey technology that combines both an SRAM-PUF and a Bus-Keeper-PUF, along with state-of-the-art security enhancements unequaled by any other FPGA.

The following illustration shows the PolarFire FPGA's dual SRAM-PUF and bus-keeper-PUF being used to generate a hardware intrinsic key and a true random number.

Figure 15 Dual SRAM-PUF and Bus-Keeper-PUF

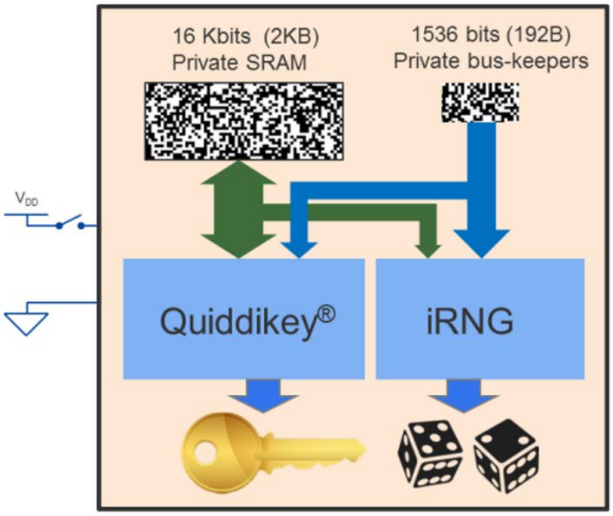

A physically unclonable function exploits intrinsic device-to-device differences generated randomly during manufacturing to create a unique secret ID, or "fingerprint," to derive a repeatable AES-256 key-encryption-key used for wrapping and storing other keys. Along with other countermeasures, the SRAM and the bus-keepers are powered-down when not in use, keeping the PUF secret ID and any keys protected by it secure.

Data Security Features

Many of the same features that provide a solid hardware-based root-of-trust for design security also make PolarFire FPGAs ideal for data security applications.

Full-Featured Crypto-Processor

PolarFire FPGAs include a TeraFire® EXP-F5200B crypto-processor dedicated to the FPGA user. The TeraFire core implements many of the most commonly used cryptographic algorithms such as AES, SHA 2, ECC, RSA, DH, and includes a cryptographic-grade TRNG.

The performance of the user's TeraFire crypto-processor should be suitable for many applications, reducing the costs (area, power, and licensing-related) compared with adding an accelerator to the FPGA fabric.

Table 1 Expected TeraFire® EXP-F5200B Crypto-processor Throughput

TeraFire EXP-F5200B Protocols and Expected Throughput[1]	
DRBG	120 Mbps[2]
AES-256	245 Mbps[2]
SHA-256	145 Mbps[2]
ECDSA-256	27/12 ms[3]
DSA-2048	64/57 ms[3]
SSA-2048	135/4 ms[3]

[1] With DPA countermeasures "on"

[2] Average over a long message

[3] Single SigGen/SigVer execution, respectively

Table 2 TeraFire® EXP-5200B Crypto-Processor Supported Algorithms and Certifications

TeraFire EXP-F5200B Supported Protocols/Features
TRNG: SP800-90A CTR_DRBG-256; SP800-90B (draft) NRBG
AES-128/192/256 E/D (ECB, CBC, CTR, OFB, CFB, GCM, KeyWrap)
SHA-1/224/256/384/512
HMAC-SHA-256/384/512; GMAC-AES; CMAC-AES
SHA-256 Key Tree
ECC: NIST P256/384/521 and Brainpool P256/384/512 curves; KeyGen, KAS - ECC CDH, ECDSA SigGen and SigVer, PKG, PKV
IFC: 1024/1536/2048/3072/4096/8192; RSA E/D; SSA_PKCS1_V1_5 SigGen & SigVer; ANSI X9.31 SigGen and SigVer
FFC: 1024/1536/2048/3072/4096; KAS - DH, DSA SigGen and SigVer

The following TeraFire EXP-F5200B NIST CAVP certifications have been awarded by NIST.

- AES: 3950, 3951
- DSA: 1077
- RSA: 2018
- ECDSA: 867, 868
- SHS: 3258, 3259
- DRBG: 1153, 1154
- HMAC: 2573

For details on which specific algorithms and modes were certified, see the NIST CAVP website using the certification numbers. Where a second certificate number is shown, it is for the TeraFire EXP-F5200ASR used by the system controller for FPGA design security, which is also certified ECC CDH: 790.

All algorithms using a secret key are available heavily protected against side-channel analyses, such as SPA, TA, DPA, and DEMA. The availability of so many common algorithms with the high level of performance and DPA resistance offered by the TeraFire crypto-processor is not only unique amongst FPGAs, but is hard to find elsewhere in any type of publicly available device.

Conclusion

Microsemi's non-volatile FPGAs have maintained a leadership position in reliability, low-power, and security for several product generations. With the end of Moore's law and Dennard scaling, along with the exponential growth rate of connect devices and human interaction with those devices, the value of our leadership in those domains has been enhanced. With the introduction of the PolarFire family of FPGAs, Microsemi has made significant advances in all domains: Microsemi has taken a leadership position with cost-optimized mid-range devices, delivering the lowest power at densities up to 500K logic elements (LEs) for communications, defense, and industrial markets, and Microsemi is now a true broad-range FPGA supplier—offering families from 1K to 500K LEs, addressing mainstream applications while delivering all the benefits of non-volatile technologies. With the introduction of the PolarFire family, the market now has a cost-optimized mid-range FPGA solution that delivers outstanding power efficiency, significantly higher security, and reliability compared to alternative solutions.

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo, CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www.microsemi.com.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.