



PolarFire® SoC Product Overview

Overview

PolarFire® SoC is built upon the award-winning PolarFire FPGA non-volatile FPGA platform. Featuring a five core Linux capable processor subsystem based on the RISC-V ISA, PolarFire SoC brings to market a royalty-free, innovative, mid-range, embedded compute platform that inherits all the benefits of the PolarFire FPGA product family. The RISC-V CPU micro-architecture implementation is a simple, 5-stage single issue in order pipeline that does not suffer from the Meltdown and Spectre exploits found in common out-of-order machines. All five CPU cores are coherent with the memory subsystem allowing a versatile mix of deterministic real time systems and Linux in a single, multi-core CPU cluster. With Secure Boot built-in, innovative Linux and Real Time modes, a large Flexible L2 memory subsystem, and a rich set of embedded peripherals, PolarFire SoC provides designers new choices in secure, power-efficient, embedded compute platforms. This document describes the features of PolarFire SoC extended commercial (0 °C to 100 °C T_j) and industrial (–40 °C to 100 °C T_j) device offerings.

Microprocessor Subsystem Features

- 64-bit RV64GC Quad Application processing cores, Fmax of 667 MHz (–40 °C to 100 °C T_j), 3.0 CoreMarks®/MHz, 2.0 DMIPs/MHz
 - L1 memory subsystem with single-error correct, double-error detect (SECDED)
 - 32 Kbytes 8-way instruction cache or optional 28 Kbytes tightly integrated memory
 - 32 Kbytes 8-way data cache
 - Memory Management Unit (MMU)
 - Physical Memory Protection (PMP) unit
- 64-bit RV64IMAC monitor processor core, Fmax of 667 MHz (–40 °C to 100 °C T_j), 3.0 CoreMarks®/MHz, 2.0 DMIPs/MHz
 - L1 memory subsystem with SECDED
 - 16 Kbytes 2-way instruction cache
 - 8 Kbytes scratch pad memory
 - PMP unit
- Flexible 2 MB L2 memory subsystem with SECDED configurable as:
 - 16-way set associative L2 cache
 - Loosely Integrated Memory (LIM) mode for deterministic access
 - Coherent Scratchpad Memory mode for shared messages across cores
- Integrated 36-bit DDR4/DDR3/LPDDR4/LPDDR3 memory controller with SECDED
 - DDR4 at 1.6 Gbps with a 8 GB address reach
- Cache coherent CPU bus matrix
- AMBA I/O switch with QoS and memory protection
- Integrated 128 Kbytes embedded non-volatile memory (eNVM) for boot
- Boot options
 - Microchip secure boot
 - User defined, PUF-protected secure boot
 - Boot directly from eNVM

-
-
- Platform interrupt controller
 - 185 interrupt sources from the microprocessor subsystem and FPGA fabric with seven priority levels
 - Local interrupt controller
 - 48 local interrupts sourced from the FPGA drive the local interrupt controller on each core
 - Debug
 - Ten hardware triggers per CPU (triggers can be configured as a breakpoint or a watchpoint)
 - Instruction trace on all CPUs
 - Performance counters
 - Runtime-configurable AXI bus monitors
 - Monitor AXI commands to DDR
 - Monitor an AXI port going into or out of the AMBA I/O AXI switch
 - 32-bit fabric monitor
 - SmartDebug
 - Dynamically monitor any two nets in the FPGA on two pins without changing the FPGA design
 - Read/write to FPGA flip-flops and memories
 - Halt clock trees, inspect logic tree
 - FPGA breakpoints
 - SmartDebug integrated into processor debug transport layer—debug from a single tool chain
 - Secure debug remotely over Ethernet (both the processor subsystem and the FPGA design)
 - Processor I/O
 - Two GigE MACs
 - An USB 2.0 OTG
 - MMC 5.1 SD/SDIO
 - Two CAN 2.0 A and B
 - Execute in place Quad SPI flash controller
 - Five multi-mode UARTs
 - Two SPI, two I²C
 - RTC, GPIO
 - Five watchdog timers
 - Timers
 - Processor to FPGA Interconnect
 - Two 64-bit AXI4 processor-to-fabric interfaces
 - Three 64-bit AXI4 fabric-to-processor interfaces
 - A 32-bit APB processor-to-fabric interface

FPGA Features

- Up to 461K logic elements consisting of a 4-input look-up table (LUT) with a fractureable D-type flip-flop
- 20 Kbytes dual- or two-port large static random access memory (LSRAM) block with built-in SECEDED
- 64 × 12 two-port μ RAM block implemented as an array of latches
- 18 × 18 math block with a pre-adder, a 48-bit accumulator, and an optional 16-deep × 18 coefficient ROM
- Built-in μ PROM, modifiable at program time and readable at run time for user data storage
- High-speed serial connectivity with built-in, multi-gigabit, multi-protocol transceivers from 250 Mbps to 12.7 Gbps
- Integrated dual x4 PCIe Gen2 endpoint (EP) and root port (RP) designs
- High-speed I/O (HSIO) supporting up to 1600 Mbps DDR4, 1333 Mbps DDR3L, and 1333 Mbps LPDDR3/DDR3 memories with integrated I/O digital
- General-purpose I/O (GPIO) supporting 3.3 V, built-in CDR for serial gigabit Ethernet, 1067 Mbps DDR3, and 1250 Mbps LVDS I/O speed with integrated I/O digital logic
- Low-power, phase-locked loops (PLLs) and delay-locked loops (DLLs) for high precision and low jitter
- 1.0 V and 1.05 V operating modes

Low-Power Features

- Low device static power
- Low inrush current
- Low-power transceivers

Reliability Features

- FPGA configuration cells single-event upset (SEU) immune
- Built-in SECDDED and memory interleaving on FPGA fabric LSRAMs
- SECDDED on all processor memories
 - Error signals trapped and exported to the FPGA fabric
- System controller suspend mode for safety-critical designs

Security Features

- Cryptography Research Incorporated (CRI)-patented differential power analysis (DPA) bitstream protection
- Integrated dual physically unclonable function (PUF)
- 56 Kbytes of secure, non-volatile memory (sNVM)
- Built-in tamper detectors and countermeasures
- Digest integrity check for FPGA, μ PROM, sNVM, and eNVM

SoftConsole Embedded IDE

- Eclipse IDE
- Firmware catalog for device drivers

Antmicro Renode™

- Open Source PolarFire SoC System Modeling environment integrated with SoftConsole

Libero® SoC PolarFire FPGA Toolset

- Complete FPGA development environment
- Includes Synplify Pro synthesis and Mentor ModelSim ME simulation

Table of Contents

Overview.....	1
1. Microprocessor Subsystem Features.....	1
2. FPGA Features.....	2
1. Block Diagram.....	6
2. Product Family Table.....	7
3. Microprocessor Subsystem.....	9
3.1. CPUs.....	9
3.2. Debug.....	13
3.3. Interrupts.....	14
3.4. Memory Subsystems.....	15
3.5. Processor I/O.....	17
3.6. Processor-to-Fabric Interconnect.....	22
3.7. Secure Boot.....	22
3.8. Peripheral Memory SECEDED Reporting and Error Injection.....	22
3.9. DMA Controller.....	23
4. Programmable Logic Subsystem.....	24
4.1. Clock Management.....	24
4.2. Debug Probe System.....	25
4.3. I/Os.....	25
4.4. Non-Volatile FPGA Fabric.....	32
4.5. PCI Express.....	35
5. System Controller	38
5.1. System Services.....	38
5.2. Programming.....	38
6. Low Power.....	40
6.1. Non-Volatile Technology.....	40
6.2. Low-Power Transceiver Lane.....	40
7. Reliability.....	41
7.1. FPGA Fabric.....	41
7.2. LSRAM.....	41
7.3. µSRAM.....	41
7.4. Digests.....	41
7.5. System Controller Suspend Mode.....	41
8. Security.....	43
8.1. Design Security.....	43
9. PolarFire SoC Device Offerings.....	44
10. Ordering Information.....	45
11. Revision History.....	46

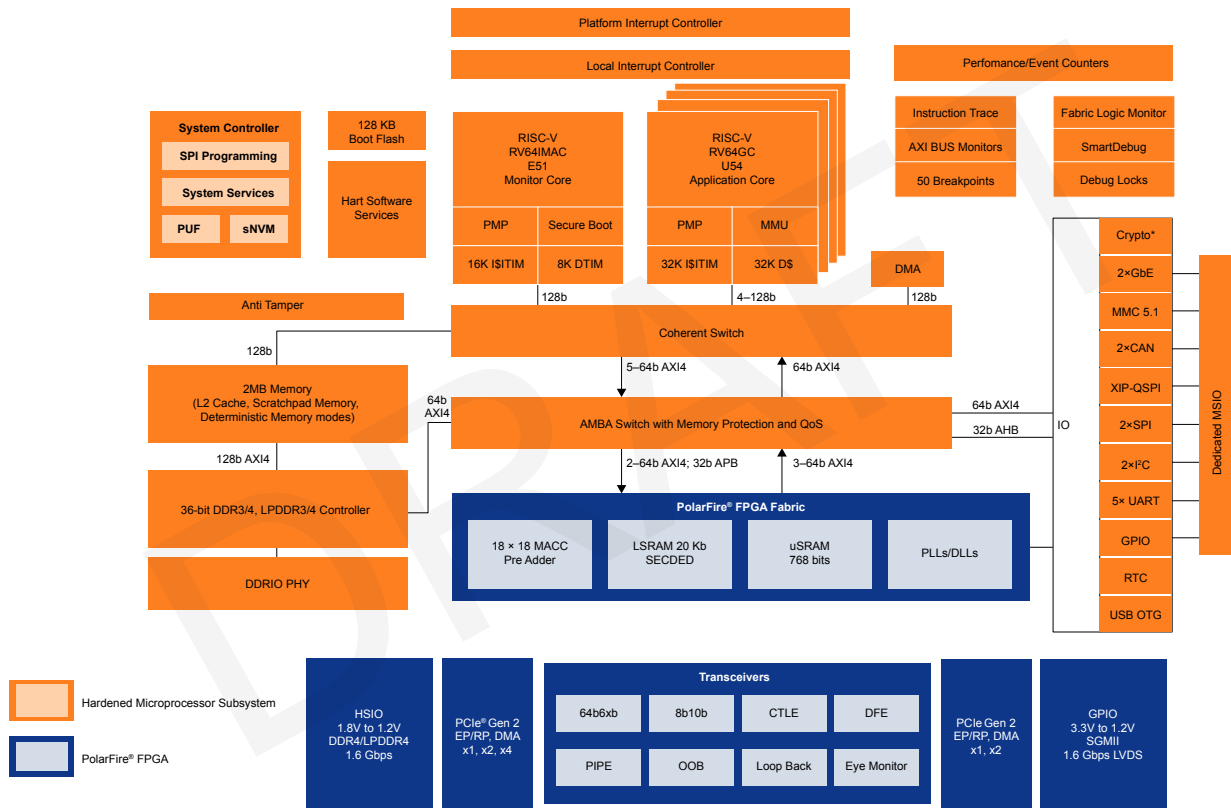
The Microchip Website.....	47
Product Change Notification Service.....	47
Customer Support.....	47
Microchip Devices Code Protection Feature.....	47
Legal Notice.....	48
Trademarks.....	48
Quality Management System.....	49
Worldwide Sales and Service.....	50

1. Block Diagram

The following illustration shows the functional blocks of PolarFire SoC.

Figure 1-1. Block Diagram

PolarFire® SoC Block Diagram



2. Product Family Table

The following table lists the product overview and packaging overview of the PolarFire SoC product family. The processor subsystem is common to all devices in the product family.

Table 2-1. PolarFire SoC Product Family

	Features	MPFS025T	MPFS095T	MPFS160T	MPFS250T	MPFS460T
FPGA fabric	K Logic Elements (4LUT + DFF)	23	93	161	254	461
	Math Blocks (18x18 MACC)	68	292	498	784	1420
	LSRAM Blocks (20 kb)	84	308	520	812	1460
	uSRAM Blocks (64x12)	204	876	1494	2352	4260
	Total RAM Mb	1.8	6.7	11.3	17.6	31.6
	uPROM Kbytes	194	387	415	470	553
	User DLLs/ PLLs	8 each	8 each	8 each	8 each	8 each
High-speed I/O	250 Mbps to 12.5 Gbps SERDES Lanes	4	4	8	16	20
	PCIe Gen2 End Points/ Root Ports	2	2	2	2	2
Total FPGA I/O	HSIO+GPIO	108	276	312	372	468
Total MSS I/O	MSS I/O	136	136	136	136	136
MSS DDR DB	MSS DDR Data Bus	16	32	32	32	32

Product Family Table

.....continued						
	Features	MPFS025T	MPFS095T	MPFS160T	MPFS250T	MPFS460T
Packaging	Type/Size/ Pitch	MSS IO/HSIO/GPIO/XCVRs				
	FCSG325 (11x11, 11x14.5*, 0.5 mm)	102/32/48/2	102/32/48/2		—	—
	FCSG536 (16x16, 0.5 mm)	—	136/60/108/4	136/60/108/4	136/60/108/4	—
	FCVG484 (19x19, 0.8 mm)	136/60/48/4	136/60/48/4	136/60/84/4	136/60/84/4	—
	FCVG784 (23x23, 0.8 mm)	—	136/144/132/ 4	136/144/168/ 8	136/144/180/8	—
	FCG1152 (35x35, 1.0 mm)	—	—	—	136/144/228/1 6	136/180/228/2 0

Notes:

1. Devices in the same package type are pin compatible.
2. Extended Commercial and Industrial temperature grade devices are available in Green RoHS packages.

3. Microprocessor Subsystem

3.1 CPUs

3.1.1 E51 Monitor Core

The E51 monitor core is a 64-bit embedded RISC-V microcontroller, including an instruction fetch unit, an execution pipeline, and a data memory system. The monitor core supports the standard RISC-V RV64IMAC user-level instruction set, with machine and user privilege modes. The E51 is primarily responsible for booting and configuring the processor subsystem after which a super loop is executed that responds to firmware service requests from the application processors.

Table 3-1. E51 Monitor Core

Feature	Description
ISA	RV64IMAC
Instruction cache	16 Kbytes two way
Data tightly integrated memory	8 Kbytes
ECC support	Single-error correct, double-error detect on the DTIM
Privileged modes	Machine (M), User (U)

3.1.1.1 E51 Instruction Fetch Unit

The E51 instruction fetch unit consists of a two-way, 16 Kbytes instruction cache that supports 64-byte cache lines. The access latency is one clock cycle. The instruction memory system is not coherent with the data memory system. Writes to memory may be synchronized with the instruction fetch stream with a FENCE.I instruction. The branch predictor comprises a branch target buffer (BTB), which predicts the target of taken branches and jumps; a branch history table (BHT), which predicts the direction of conditional branches; and a return-address stack (RAS), which predicts the target of procedure returns. The BTB is configured to hold 40 entries. The RAS is configured to hold two entries. The BHT uses a gshare prediction scheme with 7 bits of global history to access an array of 128, two-bit saturating counters. The branch predictor has a one-cycle latency, so that correctly predicted control-flow instructions result in no penalty. The branch predictor can be turned off during device configuration to create deterministic systems.

3.1.1.2 E51 I-Cache Reconfiguration

The instruction cache can be partially reconfigured into an Instruction Tightly Integrated Memory (ITIM), which occupies a fixed address range in the memory map. ITIM provides high performance, predictable instruction delivery. Fetching an instruction from ITIM is as fast as an instruction-cache hit, with no possibility of a cache miss. ITIM can hold data as well as instructions, though loads and stores to ITIM are not as performant as loads and stores to DTIM. The instruction cache can be configured as ITIM for all ways except for 1 in units of cache lines (64 bytes). A single instruction cache way must remain as an instruction cache. ITIM is allocated simply by storing to it. A store to the n^{th} byte of the ITIM memory map reallocates the first $n+1$ bytes of instruction cache as ITIM, rounded up to the next cache line. ITIM is deallocated by storing zero to the first byte after the ITIM region. The deallocated ITIM space is automatically returned to the instruction cache. For determinism, software must clear the contents of ITIM after allocating it. It is unpredictable whether ITIM contents are preserved between deallocation and allocation.

3.1.1.3 E51 Execution Pipeline

The E51 execution unit is a single-issue, in-order pipeline. The pipeline comprises five stages: instruction fetch, instruction decode and register fetch, execute, data memory access, and register writeback. The pipeline has a peak execution rate of one instruction per clock cycle. It is fully bypassed, so that most instructions have an apparent one-cycle result latency. There are several exceptions:

- LD and LW have a two-cycle result latency, assuming a cache hit.

- LH, LHU, LB, and LBU have a three-cycle result latency, assuming a cache hit.
- MUL, MULW, MULH, MULHU, MULHSU, DIV, DIVU, REM, REMU, DIVW, DIVUW, REMW, and REMUW have between a 2-cycle and 66-cycle result latency, depending on operand values.
- CSR reads have a three-cycle result latency.

The pipeline only interlocks on read-after-write and write-after-write hazards, so instructions may be scheduled to avoid stalls.

The iterative multiplier is configured to produce 16 bits per cycle with an early-out option. The iterative divider has latency of between three and 66 cycles and an early-out option.

Branch and jump instructions transfer control from the memory access pipeline stage. Correctly predicted branches and jumps incur no penalty, whereas mispredicted branches and jumps incur a three-cycle penalty. Most CSR writes result in a pipeline flush, a five-cycle penalty.

3.1.1.4 E51 Data Memory System

The E51 data memory system consists of 8 Kbytes data tightly-integrated RAM (DTIM). The access latency is two clock cycles for full words and three clock cycles for smaller quantities. Misaligned accesses are not supported in hardware and result in a trap to support software emulation. Stores are pipelined and commit on cycles where the data memory system is otherwise idle. Loads to addresses currently in the store pipeline result in a five-cycle penalty.

3.1.1.5 E51 Memory Protection

The E51 DTIM implements single-error correcting, double-error detecting (SECDEC) error correcting code (ECC). The granularity at which this protection is applied (the codeword) is 32 bits (with an ECC overhead of 7 bits per codeword).

3.1.1.5.1 E51 Memory Single-Bit Errors

In the case of a single-bit error in the L1 instruction cache, the error is corrected and the cache line is flushed. When a single-bit error is detected in the ITIM, the error is corrected and written back to the SRAM. When the L1 data cache encounters a single-bit error, the data cache will correct the error, invalidate the cache line, and write the line back to the next level of memory hierarchy.

3.1.1.5.2 E51 Memory Error Reporting

ECC events are reported by the Bus-Error Unit (BEU) for a given core. The BEU can be configured to generate interrupts either globally to the Platform Level Interrupt Controller (PLIC), or locally to the specific part where the ECC event occurred. When BEU interrupts are enabled, software can then be used to monitor and count ECC events. In order to detect uncorrectable ECC errors in the L1 memory system, interrupts must be enabled in the BEU. Specifically, to halt execution of a core when an uncorrectable instruction is detected, the BEU must be configured to generate a local interrupt. Uncorrectable ECC errors in the L1 system are also reported to the FPGA fabric as a HALT_CPU_n signal, where n = 0 –5 and n = 0 = the E51 monitor core.

3.1.1.6 E51 Local Interrupts

The E51 supports up to 48 local interrupt sources routed directly to the core. The local interrupts are sourced from the FPGA fabric. The E51 core receives the same 48 interrupt sources from the FPGA fabric as do the U54 cores.

3.1.2 U54 Application Cores

The U54 application core is 64-bit embedded RISC-V microprocessor, including an instruction fetch unit, an execution pipeline, and a data memory system. The monitor core supports the standard RISC-V RV64IMAFDC (RV64GC) user-level instruction set, with machine, supervisor, and user privilege modes. The U54s are primarily responsible for running a rich operating system such as Linux.

Table 3-2. U54 Application Cores

Feature	Description
ISA	RV64IMAFDC (RV64GC)
Instruction cache	32 Kbytes, 8-way

.....continued	
Feature	Description
Instruction tightly integrated memory (ITIM)	Maximum of 28 Kbytes
Data cache	32 Kbytes, 8-way
ECC support	Single-error correct, double-error detect on the instruction cache/ITIM and data cache
Virtual memory support	The U54 has support for Sv39 virtual memory support with a 39-bit virtual address space, 38-bit physical address space, and 32-entry TLB
Privileged modes	Machine (M), Supervisor (S), User (U)

3.1.2.1 U54 Instruction Memory System

The instruction memory system consists of a dedicated 32 Kbytes, 8-way, set-associative, Virtually Indexed Physically Tagged (VIPT) instruction cache. The access latency of all blocks in the instruction memory system is one clock cycle. The instruction cache is not kept coherent with the rest of the platform memory system. Writes to instruction memory must be synchronized with the instruction fetch stream by executing a FENCE.I instruction. The instruction cache has a line size of 64 bytes and a cache line fill will trigger a burst access outside of the PolarFire SoC CPU core complex. The core will cache instructions from executable addresses, with the exception of the ITIM. Trying to execute an instruction from a non-executable address will result in a synchronous trap.

3.1.2.2 U54 I-Cache Reconfiguration

The instruction cache can be partially reconfigured into an Instruction Tightly Integrated Memory (ITIM), which occupies a fixed address range in the memory map. ITIM provides high-performance, predictable instruction delivery. Fetching an instruction from ITIM is as fast as an instruction-cache hit, with no possibility of a cache miss. ITIM can hold data as well as instructions, though loads and stores to ITIM are not as performant as loads and stores to DTIM. The instruction cache can be configured as ITIM for all ways except for 1 in units of cache lines (64 bytes). A single-instruction cache way must remain as an instruction cache. ITIM is allocated simply by storing to it. A store to the n^{th} byte of the ITIM memory map reallocates the first $n+1$ bytes of instruction cache as ITIM, rounded up to the next cache line. ITIM is deallocated by storing zero to the first byte after the ITIM region. The deallocated ITIM space is automatically returned to the instruction cache. For determinism, software must clear the contents of ITIM after allocating it. It is unpredictable whether ITIM contents are preserved between deallocation and allocation.

3.1.2.3 U54 Instruction Fetch Unit

The U54 instruction fetch unit contains branch prediction hardware to improve performance of the processor core. The branch predictor comprises a 40-entry branch target buffer (BTB) that predicts the target of taken branches, a 128-entry branch history table (BHT) that predicts the direction of conditional branches, and a 2-entry return-address stack (RAS) that predicts the target of procedure returns. The branch predictor has a one-cycle latency, so that correctly predicted control-flow instructions result in no penalty. Mispredicted control-flow instructions incur a three-cycle penalty. The branch predictor can be turned off during device configuration to create deterministic systems. The U54 implements the standard Compressed (C) extension to the RISC-V architecture which allows for 16-bit RISC-V instructions.

3.1.2.4 U54 Execution Pipeline

The U54 execution unit is a single-issue, in-order pipeline. The pipeline comprises five stages: instruction fetch, instruction decode and register fetch, execute, data memory access, and register writeback.

The pipeline has a peak execution rate of one instruction per clock cycle, and is fully bypassed so that most instructions have a one-cycle result latency. There are several exceptions:

- LW has a two-cycle result latency, assuming a cache hit.
- LH, LHU, LB, and LBU have a three-cycle result latency, assuming a cache hit.
- CSR reads have a three-cycle result latency.
- MUL, MULH, MULHU, and MULHSU have a 5-cycle result latency.
- DIV, DIVU, REM, and REMU have between a 2-cycle and 33-cycle result latency, depending on the operand values.

The pipeline only interlocks on read-after-write and write-after-write hazards, so instructions may be scheduled to avoid stalls.

The U54 implements the standard Multiply (M) extension to the RISC-V architecture for integer multiplication and division. The U54 has a 16-bit per cycle hardware multiply and a 4-bit per cycle hardware divide.

Branch and jump instructions transfer control from the memory access pipeline stage. Correctly predicted branches and jumps incur no penalty, whereas mispredicted branches and jumps incur a three-cycle penalty.

Most CSR writes result in a pipeline flush with a five-cycle penalty.

3.1.2.5 U54 Data Memory System

The U54 data memory system has a 8-way set-associative 32 KiB write-back data cache with 64 B cache lines. The data cache is Virtually Indexed Physically Tagged (VIPT). Access latency is two clock cycles for words and double-words, and three clock cycles for smaller quantities. Misaligned accesses are not supported in hardware and result in a trap to support software emulation. The data caches are kept coherent with a directory-based cache coherence manager, which resides in the outer L2 cache. Stores are pipelined and commit on cycles where the data memory system is otherwise idle. Loads to addresses currently in the store pipeline result in a five-cycle penalty.

3.1.2.6 U54 Atomic Operations

The U54 core supports the RISC-V standard Atomic (A) extension on regions of the memory map denoted by the attribute A. Atomic memory operations to regions that do not support them generate an access exception precisely at the core. The load-reserved and store-conditional instructions are only supported on cached regions; therefore, generate an access exception on DTIM and other uncached memory regions. See *The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Version 2.1 [1]* for more information on the instructions added by this extension.

3.1.2.7 U54 Floating Point Unit (FPU)

The U54 FPU provides full hardware support for the IEEE[®] 754-2008 floating-point standard for 32-bit, single-precision and 64-bit, double-precision arithmetic. The FPU includes a fully pipelined fused-multiply-add unit and an iterative divide and square-root unit, magnitude comparators, and float-to-integer conversion units, all with full hardware support for subnormals and all IEEE default values.

3.1.2.8 U54 Virtual Memory Support

The U54 has support for virtual memory through the use of a Memory Management Unit (MMU). The MMU supports the Bare and Sv39 modes as described in the *RISC-V Instruction Set Manual, Volume II: Privileged Architecture, Version 1.10 [2]*. The U54 MMU has a 39-bit virtual address space mapped to a 38-bit physical address space. A hardware page-table walker refills the address translation caches. Both instruction and data address translation caches are fully associative, and have 32 entries. The MMU supports 2 MB megapages and 1 GB gigapages to reduce translation overheads for large contiguous regions of virtual and physical address space. Note that the U54 does not automatically set the Accessed (A) and Dirty (D) bits in a Sv39 Page Table Entry (PTE). Instead, the U54 MMU will raise a page fault exception for a read to a page with PTE.A=0 or a write to a page with PTE.D=0.

3.1.2.9 U54 Local Interrupts

Each U54 supports up to 48 local interrupt sources that are routed directly to the core. The local interrupts are sourced from the FPGA fabric. Each U54 core receives the same 48 interrupt sources from the FPGA fabric.

3.1.2.10 U54 Memory Protection

The U54 ITIM and DTIM implement single-error correcting, double-error detecting (SECDEC) error correcting code (ECC). The granularity at which this protection is applied (the codeword) is 32-bit (with an ECC overhead of 7 bits per codeword).

3.1.2.10.1 U54 Memory Single-Bit Errors

In the case of a single-bit error in the L1 instruction cache, the error is corrected and the cache line is flushed. When a single-bit error is detected in the ITIM, the error is corrected and written back to the SRAM. When the L1 data cache encounters a single-bit error, the data cache will correct the error, invalidate the cache line, and write the line back to the next level of memory hierarchy.

3.1.2.10.2 U54 Memory Error Reporting

ECC events are reported by the Bus-Error Unit (BEU) for a given core. The BEU can be configured to generate interrupts either globally to the Platform Level Interrupt Controller (PLIC), or locally to the specific part where the ECC event occurred. When BEU interrupts are enabled, software can then be used to monitor and count ECC events. In order to detect uncorrectable ECC errors in the L1 memory system, interrupts must be enabled in the BEU.

Specifically, to halt execution of a core when an uncorrectable instruction is detected, the BEU must be configured to generate a local interrupt. Uncorrectable ECC errors in the L1 system are also reported to the F{PGA fabric as a HALT_CPU_n signal, where n = 0 –5 and n = 1 through 5 = the U54 application cores.

3.1.3 Physical Memory Protection

Each CPU in PolarFire SoC includes a physical memory protection (PMP) unit compliant with the *RISC-V Instruction Set Manual, Volume II: Privileged Architecture, Version 1.10*. The PMP unit can be used to set memory access privileges (read, write, execute) for specified memory regions. Each PMP supports 16 regions with a minimum region size of 4 bytes.

3.1.3.1 Functional Description

PolarFire SoC includes a Physical Memory Protection (PMP) unit, which can be used to restrict access to memory and isolate processes from each other. The PMP unit has 16 regions and a minimum granularity of 4 bytes. It is permitted to have overlapping regions. The PolarFire SoC PMP unit implements the architecturally defined pmpcfgX CSRs pmpcfg0 and pmpcfg2 supporting 16 regions. pmpcfg1 and pmpcfg3 are implemented but hardwired to zero.

3.1.3.2 Region Locking

The PMP allows for region locking whereby once a region is locked, further writes to the configuration and address registers are ignored. Locked PMP entries may only be unlocked with a system reset. A region may be locked by setting the L bit in the pmpicfg register. In addition to locking the PMP entry, the L bit indicates whether the R/W/X permissions are enforced on M-Mode accesses. When the L bit is set, these permissions are enforced for all privilege modes. When L bit is clear, the R/W/X permissions apply only to U-mode.

3.2 Debug

3.2.1 CPU Debug

Each CPU has up to ten breakpoint registers. Breakpoints can halt the respective CPU under the following conditions.

- Address match on Load
- Address match on Store
- Address match on Instruction Fetch
- Address match on User mode
- Address match on Supervisor mode
- Address match on Machine mode

A successful match on address can generate an exception or place the machine into debug mode.

3.2.2 Trace

PolarFire SoC includes an Instruction trace interface module. For each core, the following can be captured when an instruction is retired or trapped and trace is enabled.

- The address of the instruction
- The instruction
- Privileged mode during execution
- Trap or retired indication
- Interrupt or exception indication
- Exception cause
- Exception data

PolarFire SoC uses UltraSoC debug infrastructure to compress and transport debug information over a variety of ports.

- Ethernet
- JTAG

- FPGA fabric

3.2.3 AXI Bus Monitors

There are two AXI bus monitors within the MSS that allow, at run-time, observation of AXI transactions on the I/O switch and AXI transactions into the L2 memory subsystem. Full AXI bus observability is available on the I/O switch while the AXI transactions to the L2 memory subsystems capture all AXI signals excluding the data buses. The AXI monitors can passively filter on specific addresses, which master is performing the R/W, without affecting traffic at run time.

3.2.4 Fabric Monitor

PolarFire SoC incorporates a FPGA Fabric Monitor within the MSS. 32 General purpose inputs to the MSS from the FPGA are available for monitoring FPGA logic. Eight general purpose outputs can be driven into the FPGA fabric to control either users logic or debug instrumentation. The Fabric Monitor pipes data over the debug transport layer and out through one of the defined debug ports, most commonly Ethernet.

3.2.5 SmartDebug

Two specified user I/Os is configured (at design capture stage) as either two single-ended live probes or one differential live probe. These live probes provide read access to any register in the FPGA fabric to the output pipeline registers in the LSRAMs and to all the registers in the math block in real-time without having to re-instrument the code. A snapshot of all internal probe points is created and read out asynchronously. The live-probe feature is like a two-channel oscilloscope, whose two channels can be routed out to I/Os for external observation, and to internal ports to allow fabric design observation. Selecting different probe points within the device occurs dynamically through commands over the JTAG port using SmartDebug. Reprogramming of the device is not needed.

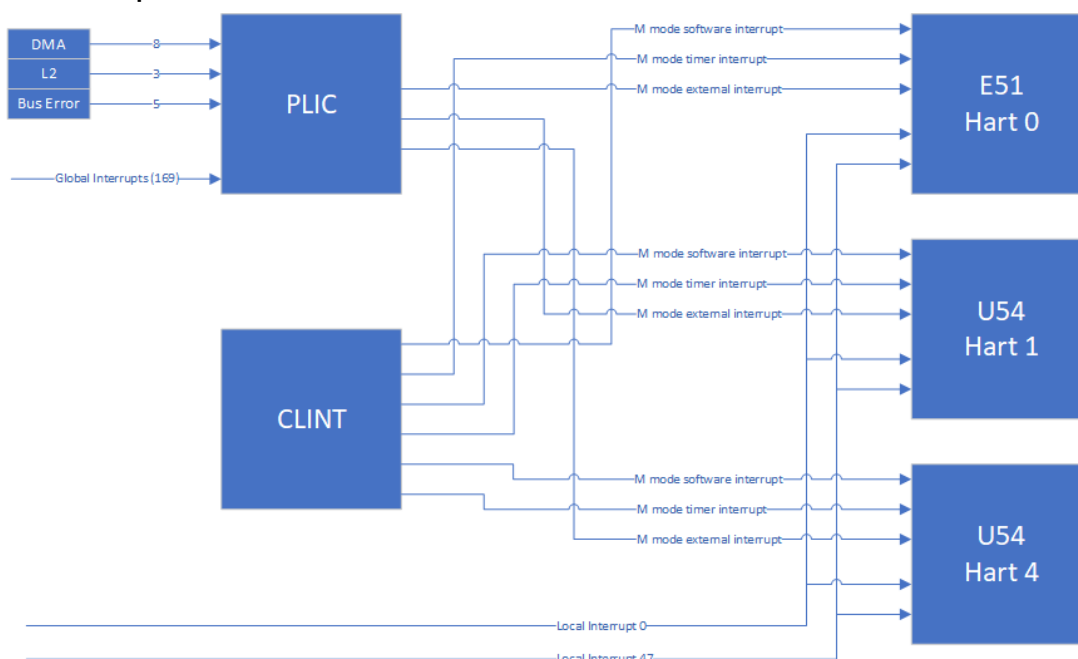
The following are included in the debug probe system.

- Active probe allows dynamic asynchronous read and write to a flip-flop or a probe point. This enables quick internal observation of the logic output or experimentation on how the logic is affected by writing to a probe point.
- Memory debug allows dynamic asynchronous read and write to a μ SRAM or a large SRAM block to quickly verify if the content of the memory is changing as expected.
- Probe insertion allows routing of nodes or debug points in the FPGA design externally through unused I/Os. An oscilloscope or logic analyzer is attached to monitor them as live signals.

3.3 Interrupts

Each hardware thread (hart) in PolarFire SoC has support for the following interrupts: local (including software and timer) and global. Local interrupts are signaled directly to an individual hart with a dedicated interrupt value. This allows for reduced interrupt latency as there is no arbitration required to determine which hart will service a given request, nor additional memory accesses required to determine the cause of the interrupt. Software and timer interrupts are local interrupts generated by the Core Local Interruptor (CLINT). Global interrupts by contrast, are routed through a Platform-Level Interrupt Controller (PLIC), which can direct interrupts to any hart in the system through the external interrupt. Decoupling global interrupts from the hart(s) allows the design of the PLIC to be tailored to the platform, permitting a broad range of attributes like the number of interrupts and the prioritization and routing schemes. By default all interrupts are handled in machine mode. The U54s, which support supervisor mode, can selectively delegate interrupts to supervisor mode.

Figure 3-1. Interrupts



3.4 Memory Subsystems

PolarFire SoC contains an on-chip 128 Kbytes embedded non-volatile memory (eNVM) for user code, a flexible L2 memory subsystem, and an integrated DDR memory controller.

3.4.1 L2 Memory Subsystem

The PolarFire SoC Level 2 Cache Controller is used to provide access to fast copies of memory for masters in a core complex. The Level 2 Cache Controller also acts as a directory-based coherency manager. The Level 2 Cache Controller offers extensive flexibility as it allows for several features in addition to the Level 2 Cache functionality, such as memory-mapped access to L2 Cache RAM for disabled cache ways, scratchpad functionality, way masking and locking, and ECC support with error tracking statistics, error injection, and interrupt signaling capabilities.

The L2 Cache Controller is configured into 4 banks where each bank contains 512 sets of 16 ways, and each way contains a 64-byte block. This subdivision into banks helps facilitate increased available bandwidth between CPU masters and the L2 Cache, as each bank has its own 128-bit inner port. As such, multiple requests to different banks may proceed in parallel. The outer port of the L2 Cache Controller is a 128-bit port shared among all banks and is connected to the DDR controller.

When cache ways are disabled, they are addressable in the L2 Loosely Integrated Memory (L2LIM) address space as described in the *UG0880 PolarFire SoC FPGA Microprocessor Sub-System User Guide*. Fetching instructions or data from the L2-LIM provides deterministic behavior equivalent to an L2 cache hit, with no possibility of a cache miss. Accesses to L2-LIM are always given priority over cache way accesses which target the same L2 cache bank. Out of reset all ways, except for way 0, are disabled. Cache ways can be enabled by writing to specific control registers. Once a cache way is enabled, it can not be disabled unless the CPU complex is reset. The highest numbered L2 cache way is mapped to the lowest L2-LIM address space, and way 1 occupying the highest L2-LIM address range. As L2 cache ways are enabled, the size of the L2-LIM address space shrinks.

The L2 Cache Controller has a dedicated scratchpad address region that allows for allocation into the cache using an address range which is not memory backed. This address region is denoted as the L2 Zero Device in the memory map. Writes to the scratchpad region will allocate into cache ways that are enabled and not masked. Care must be taken with the scratchpad, however, as there is no memory backing this address space. Cache evictions from addresses in the scratchpad will result in data loss. The main advantage of the L2 Scratchpad over the L2-LIM is that it is a cacheable region allowing for data stored to the scratchpad to also be cached in a master's L1 data cache resulting in faster access.

3.4.1.1 DDR Memory Controller

The hardened, 32-bit DDR memory controller supports the following features:

- DDR4, LPDDR4, DDR3, and LPDDR3 memory support
- Dual Rank support for dual die packages
- Max rate support 1600 Mbps
- DDR memory test feature
- Reorder queue to optimize DDR performance
- Two AXI interfaces
 - 128-bit from CPU L2 Cache
 - 64-bit from central AXI switch
- 32 outstanding transactions per AXI interface
- An integrated Clock Domain Crossing (CDC) circuit allowing the DDR controller clock to be independent of the CPU clock
- A dedicated PLL for DDR clock generation

It can support up to 8 GB of external DDR4 memory and 4 GB of DDR3 memory. SECEDED capability is also provided when configured to a 36-bit bus width, as listed in the following table.

Table 3-3. DDR Memory Controller

Configuration	Active Pads	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4 ¹
5x8 DDR with SECEDED	36	DDRx8	DDRx8	DDRx8	DDRx8	DDRx8 (4 used)
4x8 DDR	32	DDRx8	DDRx8	DDRx8	DDRx8	Not used
3x16 DDR with SECEDED	36	DDRx16		DDRx16		DDRx16 (4 used)
2x16 DDR	32	DDRx16		DDRx16		Not used
3x16 DDR with SECEDED	18	DDRx8	DDRx8	—	—	DDRx8 (2 used)
2x16 DDR	16	DDRx8	DDRx8	—	—	—
1x16 DDR with SECEDED	18	DDRx16		—	—	DDRx16 (2 used)
1x16 DDR	16	DDRx16		—	—	—
1x32 DDR	32	DDRx32				—

Note:

1. Lane 4 is only 4 bits wide, the upper data bits on the DDR memory are not connected.

3.4.1.2 Processor Interconnect

There are two interconnect switches built into the MSS. First, there is a fully populated coherent switch that manages coherence through the memory subsystems and provides a deterministic data path to the L2 memory subsystem when it is configured as a loosely integrated memory. Additionally, a central AMBA I/O switch manages the interconnect between the CPU complex, the peripheral I/O space, the hardened DDR memory controller, and the FPGA fabric. The AMBA switch also includes Quality of Service (QoS) features. The QoS feature is essentially a 4-bit value denoting priority for the data path. The central I/O switch is partially connected and supports 15 masters with nine slaves. The AMBA switch also contains a Memory Protection Unit (MPU) scheme that mimics the Physical Memory Protection (PMP) scheme defined in the RISC-V Privileged Specification. Specifically, the bus masters listed in the following table pass through the AMBA MPU. PMP region address granularity start at 4096 and increase by powers of two. Regions can be defined to support execution, read or write. An additional lock bit can be set on each register that prevents further modification to the MPU protection scheme until the next power on reset.

Table 3-4. Processor Interconnect

Master	PMP Region Count
FIC_0	16
FIC_1	16
FIC_2	8
Crypto	4
Ethernet_0	8
Ethernet_0	8
USB	4
MMC	4
DRI	8
Trace	2

3.5 Processor I/O

3.5.1 Gigabit Ethernet MAC

PolarFire SoC contains two identical Gigabit Ethernet MACs (GEM) integrated into the MSS. Each MAC can contain a maximum frame length of 10,240 bytes that are SECDED protected. Additionally, the GEM supports a built-in packet buffer DMA. Each GEM supports the following IEEE® 802 standards:

- IEEE 802.3br Frame Pre-Emption (or Interspersing Express Traffic)
- IEEE 802.1Qci Receive (Ingress) Traffic Policing
- IEEE 802.1Qbb Priority-Based Flow Control
- IEEE 802.1Q VLAN Tagging with Recognition of Incoming VLAN and Priority Tagged Frames
- IEEE 802.1AS
- IEEE 802.1Qav
- IEEE 802.1Qbv
- IEEE 1588-2002 (v1), IEEE 1588-2008 (v1 and v2)
- IEEE 802.1CB Frame Redundancy and Elimination

3.5.1.1 PHY Interfaces

Each GEM is configured to simultaneously support TBI and GMII/MII modes. When using TBI, the PCS block of the MAC is used, but not as an IEEE802.3X interface to a transceiver, but rather, is fed into a dedicated MSS SERDES block and from there, interfaces to a PHY. This serialized interface between MAC and PHY is known as SGMII and is part of the MSS. In SGMII mode, the PCS interface and the link speed auto-negotiation blocks are re-purposed from their 802.3X function and are instead used to convey control information related to the MAC-PHY interface.

3.5.1.1.1 Direct SGMII I/O

The following functionality provides an SGMII interface from the GEM to the built-in MSS SGMII PHY:

- Clock Domain Recovery (CDR) of received 125 MHz clock
- Serializing/De-serializing
- PLL for synthesis of a 125 MHz transmit clock
- I/O buffers (four I/Os per MAC instance) allowing for differential transmit and receive data pairs

Note: It is not possible to support the SyncE protocol when using the direct MSS SGMII interface.

3.5.1.1.2 GMII/MII To FPGA Fabric

A GMII/MII interface is provided between each MAC and the FPGA fabric to provide flexibility. In particular, it allows:

- Performing customized manipulation of data on-the-fly
- Connecting to a FPGA fabric transceiver channel to provide an SGMII interface. Note that in this case, the transceiver would be configured to perform 8b10b encoding/decoding in its PCS. This particular approach to implementing an SGMII interface differs from the direct hard SGMII interface of the MSS in that it allows support of the SyncE protocol, which is often used in conjunction with IEEE 1588.

3.5.1.1.3 PHY Management

Each MAC has an MDC output and an MDIO input and output port, which may be brought out separately for each MAC instance within the MSS either to MSSIO or to the FPGA fabric. If desired, however, the user could bring out only one management interface (and not use the second), as it is possible to control multiple PHYs using the one interface (if hardware separation is not required).

3.5.1.2 MSS Receive Filtering

3.5.1.2.1 Internal Filtering

The GEM is configured to have four internal specific address filters configured. Each filter can be configured to contain a MAC address, which can be specified to be compared against the source address (SA) or destination address (DA) of each received frame. There is also a mask field to allow certain bytes of the address to be excluded in the comparison. If the filtering matches for a specific frame, then it is passed on to the DMA memory. Otherwise the frame is dropped. Frames may also be filtered using the Type ID field for matching. There are four Type ID registers in the internal register space and these may be enabled individually. Hashing of the received frame's DA may be configured, as described in the *UG0880 PolarFire SoC FPGA Microprocessor Sub-System User Guide*.

3.5.1.2.2 External Filtering

To allow for more sophisticated matching of incoming frames, based on more than just addresses, for example, an external filtering interface is present. As a frame is received, the MAC parses the frame and determines what field is currently present. A strobe signal is provided to allow latching of each field in the fabric. Customized (combinational) matching circuitry can then analyze whether or not it is interested in the particular frame (filtering use case) or what receive priority queue to put the frame in (prioritizing use case).

3.5.2 MMC 5.1/SD/SDIO/eMMC

PolarFire SoC contains one MMC5.1 compliant peripheral and PHY. The controller interfaces to MSS I/O through the IOMUXs to MSSIO exclusively. PolarFire SoC MSS I/Os do not support the dynamic voltage scaling of some SD cards, external voltage level translators should be used if required. The following eMMC/SD card standards are supported in PolarFire SoC.

3.5.2.1 SD Card Standards

- Default Speed (DS)
- High Speed (HS)
- UHS-I SDR12
- UHS-I SDR25
- UHS-I SDR50
- UHS-I SDR104
- UHS-I DDR50

3.5.2.2 eMMC Standards

- Standard Speed
- High Speed
- DDR52
- HS200
- HS400
- HS400 Enhanced Strobe

3.5.3 USB 2.0 OTG

PolarFire SoC includes a USB 2.0 OTG-compliant core with an ULPI interface to the IOMUX, and then on to dedicated MSS I/O. The USB core supports the following features:

- Operates either as the function controller of a high/full-speed USB peripheral or as the host/peripheral in point-to-point or multi-point communications with other USB functions
- Complies with the USB 2.0 standard for high-speed (480 Mbps) functions and with the on-the-go supplement to the USB 2.0 specification
- Supports OTG communications with one or more high-, full-, or low-speed device
- Supports Session Request Protocol (SRP) and Host Negotiation Protocol (HNP)
- Supports suspend and resume signaling
- Supports Link Power Management
- Offers dynamic allocation of endpoints to maximize number of devices supported
- Number of Tx endpoints: 4 (plus control endpoint)
- Number of Rx endpoints: 4 (plus control endpoint)
- Multipoint capabilities supported
- Software connect/disconnect feature enabled
- High bandwidth support for ISO enabled endpoints
- Hardware-selectable option for 8-bit/4-bit LPI interface
- Vendor control and status register enabled with width of 4 and 8, respectively
- Number of DMA channels: 4
- Dynamic FIFO support allows dynamic allocation of buffer depth per enabled endpoint

3.5.4 User Crypto

PolarFire SoC embeds an Athena TeraFire F5200B side channel resistant crypto coprocessor and is available on "S" data security devices. The user crypto core can be mapped between the MSS and the FPGA fabric via a hardware based handshaking mechanism. The crypto core also supports the F5200B streaming interface directly into the fabric. Internal memories in the Athena core are SECDED protected.

3.5.4.1 TeraFire EXP-F5200B Supported Protocols Features

- TRNG SP800-90A CTR_DRBG-2565; SP800-90B (draft) NRBG
- AES 128/192/256 key lengths E/D (ECB, CBC, CTR, OFB, CFB, GCM, KeyWrap)
- SHA-1/224/256/384/512
- HMAC-SHA-1/224/256/384/512; 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 & SigVer, PKG, PKV
- IFC: 1024/1536/2048/3072/4096 RSA E/D; SSA_PKCS1_V1_5 SigGen & SigVer; ANSI X9.31 SigGen & SigVer
- FFC: 1024/153/2048/3072/4096; KAS - DH, DSA SigGen & SigVer

3.5.5 Controller Area Network

PolarFire SoC integrates two controller area network cores compliant to CAN 2.0 A and B and conforms to ISO 11898-1. Internal message SRAM is SECDED-protected. The following are CAN controller features.

3.5.5.1 Receive Path

- 32 receive buffers
- Each buffer has its own message filter
- Message filter covers: ID, IDE, RTR, data byte 1, and data byte 2
- Message buffers can be linked together to build a bigger message array
- Automatic remote transmission request (RTR) response handler with optional generation of RTR interrupt

3.5.5.2 Transmit Path

- 32 Tx message holding registers with programmable priority arbitration
- Message abort command
- Single-shot transmission (no automatic re-transmission upon error or arbitration loss)

3.5.5.3 Debugging Support

- Listen only mode
- Internal loop back mode
- External loop back mode
- Error capture register
- SRAM test mode to support software based memory testing (SRAM is addressable by the CPUs if the core is disabled)

3.5.6 QSPI XIP controller

PolarFire SoC integrates a Quad SPI (QSPI) flash controller with eExecute in place (XIP) capabilities. The following are QSPI features.

- Master SPI Data Rate
 - Programmable SPI clock HCLK/2, HCLK/4 or HCLK/6
 - Maximum data rate is HCLK/2
- FIFOs
 - Transmit and Receive FIFO
 - 16-byte transmit FIFO depth
 - 32-byte receive FIFO depth
- SPI Protocol
 - Master operation
 - Motorola SPI supported
 - Slave select operation in idle cycles configurable
 - Supports extended SPI operation (1, 2, and 4 bits)
 - Supports QSPI operation (4-bit operation)
 - Supports BPSPI operation (2-bit operation)
 - Support XIP (execute in place)
 - Supports 3 or 4-byte SPI address.
- Frame Size
 - Supports 8-bit frames directly
 - Back-to-back frame operation supports >8-bit frames
 - Supports up to 4GB Transfer (2^{32} bytes)
- Processor Overhead Reduction
 - Support of SPI flash command/data packets with automatic data generation and discard function

3.5.7 Serial Peripheral Interface

Serial peripheral interface (SPI) is a synchronous serial data protocol that enables the microprocessor or microcontroller and peripheral devices to communicate with each other. The SPI controller provides a serial interface compliant with the Motorola SPI, Texas Instruments synchronous serial, and National Semiconductor MICROWIRE™ formats. In addition, SPI supports interfacing with large SPI flash and EEPROM devices and a hardware-based slave protocol engine. There are two identical SPI controllers in the MSS.

- The SPI peripherals support the following features.
- Master and Slave modes
- Selectable slaves (up to 8)
- Configurable slave select operation

- Configurable clock polarity
- Separate transmit (Tx) and receive (Rx) FIFOs to reduce interrupt service loading

3.5.8 Multi-Mode UART

The PolarFire SoC multi-mode universal asynchronous/synchronous receiver/transmitter (MMUART) performs serial-to parallel conversion on data originating from modems or other serial devices, and performs parallel-to serial conversion on data from the CPUs. PolarFire SOC contains 5 identical MMUARTs. The MMUART is software compatible with the popular 16550 UART device.

The MMUART peripherals support the following features.

- Asynchronous and synchronous operations
- Full programmable serial interface characteristics
- Data width is programmable to 5, 6, 7, or 8 bits
- Even, odd, or no-parity bit generation/detection
- 1, 1½, and 2 stop bit generation
- 9-bit address flag capability used for multi-drop addressing topologies
- Separate transmit (Tx) and receive (Rx) FIFOs to reduce processor interrupt service loading
- Single-wire half-duplex mode in which Tx pad can be used for bi-directional data transfer
- Local interconnect network (LIN) header detection and auto baud rate calculation
- Communication with ISO 7816 smart cards
- Fractional baud rate capability
- Return to Zero Inverted (RZI) mod/demod blocks that allow infrared data association (IrDA) and serial infrared (SIR) communications
- The most significant bit (MSB) or the least significant bit (LSB) as the first bit while sending or receiving data

3.5.9 I²C

Philips inter-integrated circuit (I²C) is a two-wire serial bus interface that provides data transfer between many devices. PolarFire SoC contains two identical I²C peripherals in the MSS that provide a mechanism for serial communication between the PolarFire SoC and external I²C compliant devices. I²C peripherals support the following features:

- Master and Slave modes
- 7-bit addressing format and data transfers up to 100 Kbps in Standard mode and up to 400 Kbps in Fast mode
- Multi-master collision detection and arbitration
- Own slave address and general call address detection
- Second slave address detection
- System management bus (SMBus) timeout and real-time idle condition counters
- Input glitch or spike filters

3.5.10 Real Time Counter

PolarFire SoC integrates a real time calendar with built-in clock prescaler and an alarm wake-up comparator. It has the following two modes of operation.

- Real time Calendar (counts seconds, minutes, hours, days, weeks, months, and years)
- Binary counter (counts from 0 to 2⁴³)

3.5.11 Watchdog Timer

There are five watchdogs in the MSS, one per CPU. The watchdog guards against system crashes by requiring that it is regularly serviced by the assigned processor. The watchdog is not enabled at reset and generates a NMI upon reaching the trigger value. Once enabled, the watchdog cannot be disabled.

3.5.12 Timer

The PolarFire SoC system timer consists of two programmable 32-bit decrement counters that generate interrupts to the Cortex-M3 processor and FPGA fabric. The two 32-bit timers are identical and have the following features.

- One-shot mode
- Periodic mode
- Concatenation mode in which two 32-bit timers can be concatenated to create a 64-bit timer
- Option to enable or disable the interrupt requests when timer reaches zero
- Controls to start, stop, and reset the timer

3.6 Processor-to-Fabric Interconnect

The interconnect between the FPGA fabric and the MSS switches are either AXI- or APB-based. There are three AXI 64-bit interfaces. Two interfaces support a full, 64-bit AXI bus into the fabric from the MSS and from the fabric to the MSS. One AXI 64-bit interface supports a full, 64-bit AXI bus from the fabric into the MSS. The AXI buses all have DLLs to cancel out any insertion delay and can operate asynchronously to the processors operating frequency. In addition, a 32-bit APB bus is provided to the fabric from the MSS.

Table 3-5. Processor-to-Fabric Interconnect

Interface	Type	Width	MSS to Fabric	Fabric to MSS	DLL
FIC_0	AXI4	64b	Yes	Yes	Yes
FIC_1	AXI4	64b	Yes	Yes	Yes
FIC_2	AXI4	64b	No	Yes	Yes
FIC_3	APB	32b	Yes	No	Yes

3.7 Secure Boot

PolarFire SoC comes with two secure boot options. For the default PolarFire SoC secure boot method, the system controller will copy the Microchip secure boot loader from its private, secure memory area and load it into the 8 Kbytes DTIM of the E51 monitor core. Reset will be released to the CPUs and the boot code will start executing. The default secure boot loader will perform a signature check on the 128 Kbytes eNVM then run a hash on the eNVM image. If no errors are reported, the code will jump to the eNVM. If errors are reported, the system controller will activate a tamper alarm that asserts a signal to the FPGA fabric. Users can then decide on a plan of action.

The second secure boot method allows users to place their own boot code in the secure non-volatile memory (sNVM) area of the chip. The sNVM is a 56 Kbytes nonvolatile memory that can be protected by the built-in Physically Unclonable Function (PUF), meaning the unique PUF ID can serve as an initialization vector for an AES encrypt/decrypt operation performed by the side-channel resistant system controller co-processor. On power-up, the system controller will copy the user code from sNVM and write it to the E51 monitor core DTIM. From there, your custom secure boot loader starts executing.

3.8 Peripheral Memory SECDED Reporting and Error Injection

The Gigabit Ethernet MACs, the MMC 5.1 controller, the USB OTG controller, the CAN controllers, the crypto core and memory in the built-in MSS DDR controller are protected with by single-error correct, double-error detect (SECDED) error correction code (ECC) subsystem, which adds 7 bits to 32-bit memories and 8 bits to 64-bit memories. The memories within the CPU system have their own ECC control and reporting systems. The external DDR memory supports optional ECC (using a fifth DDR bank). Each MSS internal memory system has its own set of control and status registers, which consists of a status, interrupt enable, count, and error injection registers. When a two-bit error is detected, the data will not be corrected. If the data is being read over an internal AMBA bus, the system will respond with an APB, AHB, or AXI error response marking the data as corrupted. If the interrupt is enabled, an interrupt will also be generated. If the data is not being read over an AMBA bus like USB, CAN, Ethernet transmit data then corrupted data will be transmitted, and an interrupt generated. The system can then respond to the bus error event or interrupt and take the appropriate recovery action. ECC error injection is supported to ease

customer validation of the error correcting subsystem. Data may be written with 1-, 2-, or 3-bit errors by setting the appropriate EDAC error injection control registers.

3.9 DMA Controller

The PolarFire SoC direct memory Access (DMA) controller supports up to 4 channels of independent simultaneous transfers. Each channel has its own set of controller registers and two interrupts, complete and error. Bus transaction sizes are programmable and the transactions can be auto-loaded into the DMA engine. The DMA engine works in conjunction with hart system services (firmware running on the E51).

4. Programmable Logic Subsystem

The following section describes the programmable logic subsystem.

4.1 Clock Management

In each PolarFire FPGA, there are eight DLLs and eight PLLs to provide flexible clock generation and management capabilities. In addition to these DLLs and PLLs, up to 15 transceiver lane transmit PLLs are also available.

The following are key highlights of the clock management architecture.

- High-speed buffers and routing for low-skew clock distribution
- Frequency synthesis and phase shifting
- Low-jitter clock generation and jitter filtering

4.1.1 DLL

The DLL provides a calculated PVT compensated delay to the I/O's digital delay lines and delay or phase-shifted clocks to the FPGA fabric.

The following are the major modes to which the DLL can be configured.

- Time reference mode—the DLL takes a single clock as an input and determines how many delay line buffer taps are required for a signal to pass through them to rotate a signal. The main use of time reference mode is to know how many delay taps are needed to delay the clock by 90 degrees. The value is then provided to the data strobe signal (DQS)/DQSn input signals for double data rate (DDR) memory controllers to delay all DQS/DQSn signals by the required 90-degree phase shift to capture the data from the memory devices. Multiple memory interfaces of the same clock rate can reuse the same DLL with lane level controls for PVT updates.
- Clock injection delay mode—the DLL can be used to compensate for the clock injection delay associated with the source synchronous receive interfaces. The DLL can match delays for the global, regional, and high-speed bank clocks. There are two outputs from the DLL in this mode: a x1 output fixed in time and another output that can be divided by x1, x2, or x4 and can be phase shifted.

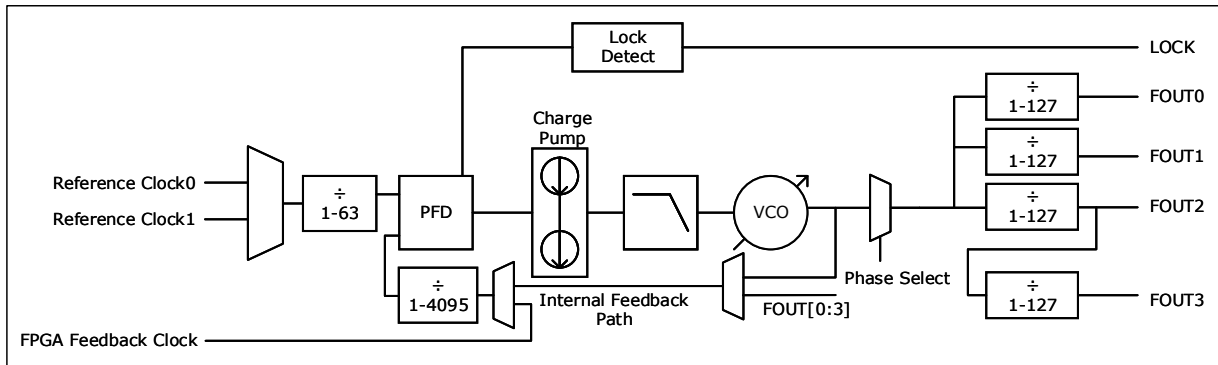
4.1.2 PLL

The programmable delta-sigma, low-jitter fractional PLLs are multi-function and general-purpose frequency synthesizers, as shown in the [PLL Block Diagram](#). Wide input and output ranges along with the best-in-class jitter performance allow these PLLs to be used for almost any clocking application. With excellent supply noise immunity, the PLL is ideal for use in noisy FPGA environments.

- The PLL output clock is available in eight phases with 45-degree phase differences. All eight phases are selectable to drive four separate outputs from the PLL, where each output can select any of the eight phases independent of other output selections and each output can also be driven to a zero output when not used.
- Each of the four outputs from the PLL can then be divided independently for any value from 1 to 127. Each of the PLL outputs can have the output divider released by up to seven VCO/4 cycles. The delayed outputs can be set independently for each output clock.
- Fractional-N (24-bit accuracy) capability is added to the feedback divider to have the VCO frequency be a non-integer divide of the reference clock input frequency. The base frequency is applied to all PLL outputs.
- The PLL supports glitch-free start and stop on any one of the four outputs independently by either a register map or a fabric control. This capability also allows the output divider values and the VCO/4 phase selection to be modified glitch-free during the time that the clock is stopped.
- For fine granularity phase control of the PLLs, they can be cascaded with DLLs located near the PLLs, whereby the DLL delay lines can be used in a process, voltage, and temperature (PVT) compensated or non-PVT compensated mode to provide the phase control needed.

The following illustration shows the flow of the PLL functionality.

Figure 4-1. PLL Block Diagram



4.1.3 Clock Network

The clock network is designed to route clocks and asynchronous reset signals to large sections of the fabric with limited skew. On occasion, the network can also be used for other high fanout signals that can tolerate long delays, such as non-timing-critical synchronous enables or resets. There are two main clock networks for the FPGA fabric, global, and regional clocks.

4.1.3.1 Global Clocks

There are 24 clocks on the device with global, low-skew scope to all synchronous elements. The global can be divided into left and right sides of the device. Thus, the number of global clocks can increase to 48 total clocks with 24 in the left and 24 in the right.

4.1.3.2 Regional Clocks

There are up to 38 regional clock domains that interface to the edges of the device. The regional clocks provide a fixed number of logic elements based on the size of the device. Up to 14 clocks are available for the FPGA I/Os and up to 24 clocks are available for the transceiver lanes, one for each lane direction. These are the fast insertion clock networks used to move data in and out of the fabric.

4.2 Debug Probe System

Two specified user I/Os can be configured (at design capture stage) as either two, single-ended live probes or one, differential live probe. These live probes can provide read access to any register in the FPGA fabric, to the output pipeline registers in the LSRAMs, and to all the registers in the math block in real time without having to re-instrument the code. A snapshot of all internal probe points can be created and read-out asynchronously. The live-probe feature can be considered a two-channel oscilloscope, whose two channels can be routed out to I/Os for external observation and to internal ports for fabric design observation. Selecting different probe points within the PolarFire FPGA occurs dynamically through commands over the JTAG port using SmartDebug. Reprogramming of the FPGA is not required.

The features of the debug probe system are:

- Active probe allows dynamic asynchronous read and write to a flip-flop or a probe point. This enables quick internal observation of the logic output or experimentation on how the logic will be affected by writing to a probe point.
- Memory debug allows dynamic asynchronous read and write to a μ SRAM or a large SRAM block to quickly verify if the content of the memory is changing as expected.
- Probe insertion allows routing of nodes or debug points in the FPGA design externally through unused I/Os. An oscilloscope/logic analyzer can be attached to monitor them as live signals.

4.3 I/Os

PolarFire device user I/Os support multiple I/O standards while providing the high bandwidth needed to maximize the internal logic capabilities of the device and achieve the required system-level performance.

4.3.1 Low-Power, High-Speed Transceiver Lane

All PolarFire FPGAs contain state-of-the-art low-power transceiver lane capabilities from speeds as low as 250 Mbps up to 12.7 Gbps. The PMA is designed to support multiple protocols (as listed in the following table) with state-of-the-art control and debug features. PCI Express Gen1 or Gen2 support is provided by a hard macro. All other protocols are implemented with a soft IP. Serial Gigabit Ethernet is also supported with GPIO 3.3 V LVDS differential pairs. A single transmit PLL can provide a high-speed clock up to four transceiver lanes.

Table 4-1. Transceiver Lane Protocol Support

Protocol	Data Rate (Gbps)	Channels Bonded
PCIe	2.5, 5	1, 2, 4
Interlaken	6.375, 12.7	1–16
10GBASE-KR	10.3125–12.7	1
SGMII/QSGMII	1.25–5	1
XAUI	3.125	4
RXAUI	3.125, 6.25	2, 3, 4, 6
HiGig/HiGig+/HiGiII	3.75–4.065	4
Fiber Channel	0.6144–12.165	1
SRIO	1.25–6.3	1, 2, 4, 8
SATA	1.5–6	1
JESD204B	0.5–12.5	1–4
Display Port	2, 5, 8	4
SDI	0.277–11.88	1

4.3.1.1 Low-Power Transceiver Lane Features

The following are features of the low-power transceiver lane.

- Advanced low-power modes
- Programmable transmit amplitude and emphasis control
- Low-speed CDR operation with support for 270 Mbps SMPTE serial line rates
- Continuous time linear equalization (CTLE) and decision feedback equalization (DFE) for long-reach or backplane applications
- Auto-adaption at receiver equalization and integrated eye monitor feature for easy serial link tuning
- Eye monitor and/or equalization can be powered down to reduce power if not needed
- Out-of-band, electrical idle signaling capability for SAS, SATA, and PCIe
- Multiple loopback modes for test and debug
- Transmit jitter attenuation for loop timing applications (SyncE compatible)
- Hot-socketing capable
- IEEE 1149.6 AC JTAG
- Adjacent channel loopback modes allow transceiver lane data streams to remain active during FPGA fabric programming

4.3.1.2 Transmitter

The transmitter is fundamentally a parallel-to-serial converter with a conversion ratio of 8, 10, 16, 20, 32, 40, 64, or 80 bits. It allows the designer to trade-off data path width for timing margin in high-performance designs. These transmitter outputs drive the PC board with a differential output signal. TX_CLK is the appropriately divided serial

data clock available to the fabric, and can be used directly to register the parallel data coming from the internal logic. The transmit parallel data has additional hardware support for the 8b/10b, 64b/66b, or 64b/67b encoding schemes to provide a sufficient number of transitions. The bit-serial output signal drives two package pins with differential signals. The output signal pair supports a wide variety of serial protocols and has programmable signal swing, as well as programmable pre- and post-emphasis to compensate for PC board losses and other interconnect characteristics. For shorter channels, the swing can be reduced to lower power consumption. Each transmit lane can be sourced by one of two transmit PLLs. Each transmit PLL can drive up to four transceiver lanes. Transmitter PLLs are state-of-the-art fractional frequency synthesizers with integrated jitter attenuation.

4.3.1.3 Receiver

The receiver is fundamentally a serial-to-parallel converter with clock recovery changing the incoming bit-serial differential signal into a parallel stream of words of 8, 10, 16, 20, 32, 40, 64, or 80 bits. This allows the FPGA designer to trade off the internal data path width versus logic timing margin. The receiver takes the incoming differential data stream, feeds it through programmable linear and decision feedback equalizers (to compensate for PC board and other interconnect characteristics), and uses the reference clock input to initiate clock recognition. The data pattern uses non-return-to-zero (NRZ) encoding and optionally guarantees sufficient data transitions by using the selected encoding scheme. The outgoing parallel data has additional hardware support for the 8b/10b, 64b/66b, or 64b/67b encoding schemes to provide a sufficient number of transitions. Parallel data is transferred into the FPGA logic using the recovered clock (RX_CLK).

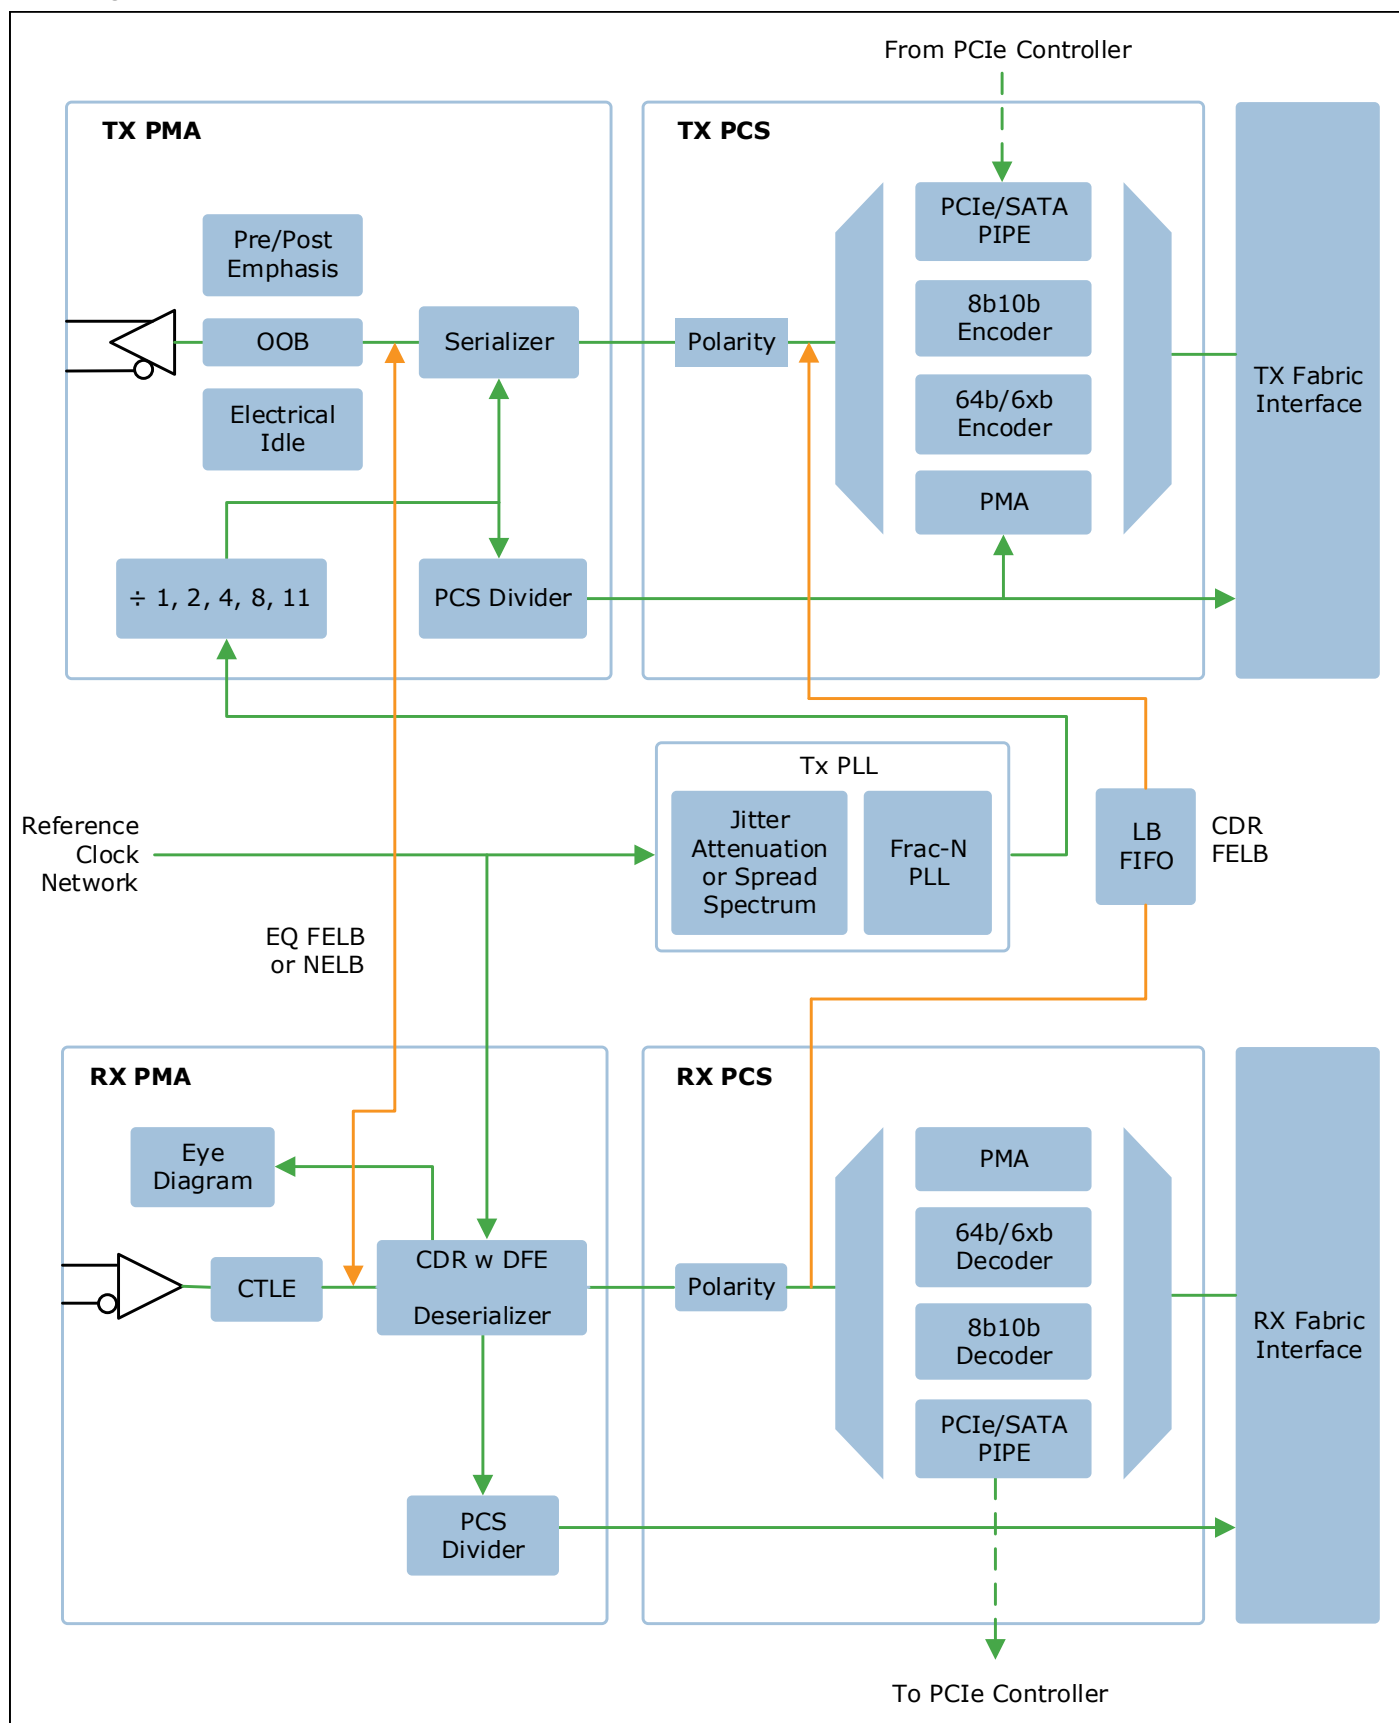
4.3.1.4 Transceiver Lane Modes

The transceiver lane supports the following five modes of operation.

- PMA—direct access to the PMA without any encoding
- 8b/10b—8b/10b encoding/decoding is provided
- 64b/6xb—64b/66b or 64b/67b encoding/decoding with gearbox logic is provided
- PIPE—a PIPE interface supporting both PCIe Gen2 and SATA 3.0
- PCIe—direct connection to the embedded PCIe Gen2 controller

The following illustration shows the collaboration of the five modes that the transceiver lanes support.

Figure 4-2. Transceiver Lane Modes

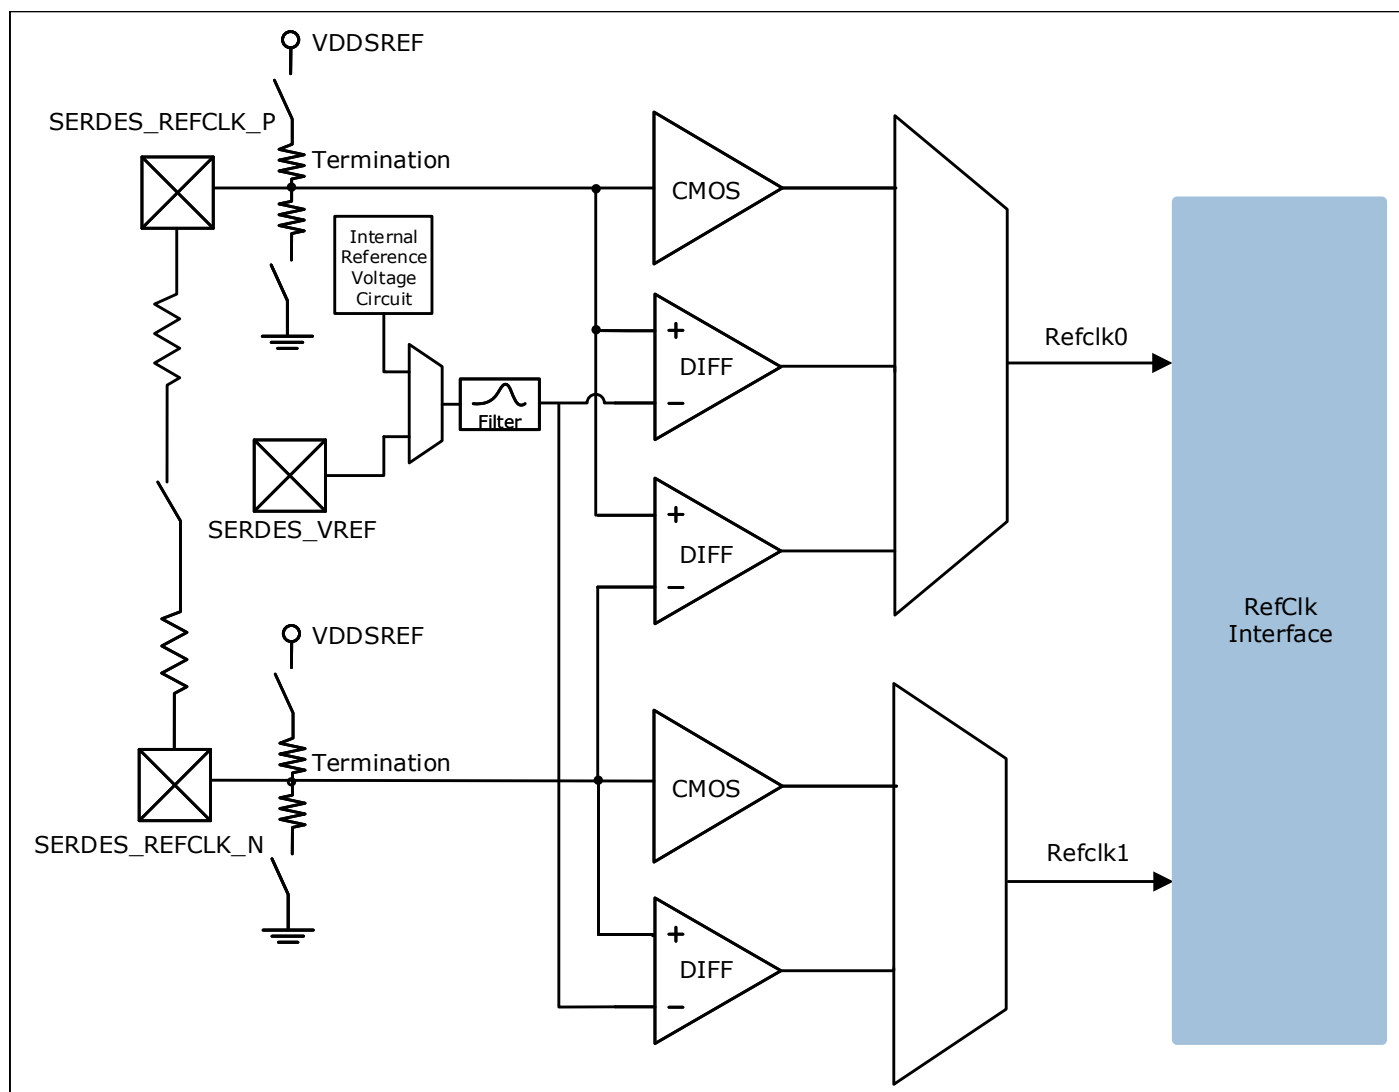


4.3.1.5 Reference Clock

The reference clock pins allow connections directly with the transceiver lane quads. The reference clock inputs provide flexibility to interface with both single-ended and differential clocks, and can drive up to two independent clocks per transceiver lane quad. These reference clocks can also be sources for the global and regional clock networks in the FPGA fabric of the device.

The following illustration shows the connectivity between the reference clock and transceiver lane quads.

Figure 4-3. Reference Clock



4.3.1.6 Quad Lane Overlay Assignments

The transceiver lane either connects the parallel side of the interface to the PCIe Gen2 controller or to the fabric. The PCIe connections are fixed in the hardware and have a dedicated number of combinations between the two controllers. The fabric interface is used to support the PMA, 8b/10b, 64b/6xb, and PIPE modes and have complete flexibility into the fabric connections.

The following table lists the combinations between the PCIe and fabric controllers.

Table 4-2. Quad0 Lane Assignments

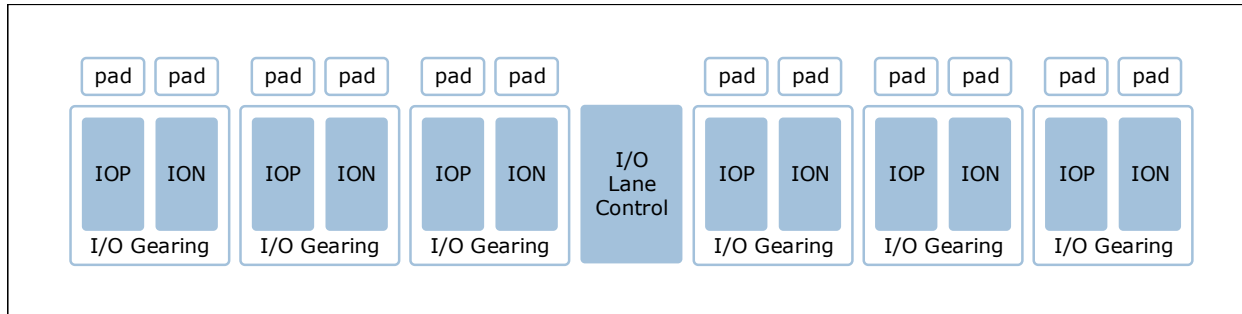
PCIe_0 Controller	Quad0 Lane 0	Quad0 Lane 1	Quad0 Lane 2	Quad0 Lane 3	PCIe_1 Controller
x1	PCIe_0	Not available	Not available	PCIe_1	x1
x1	PCIe_0	Unused	PCIe_1	PCIe_1	x2
x2	PCIe_0	PCIe_0	Not available	PCIe_1	x1
x2	PCIe_0	PCIe_0	PCIe_1	PCIe_1	x2
x4	PCIe_0	PCIe_0	PCIe_0	PCIe_0	Unused
x1	PCIe_0	Not available	Fabric	Fabric	Unused
x2	PCIe_0	PCIe_0	Fabric	Fabric	Unused
Unused	Fabric	Fabric	Not available	PCIe_1	x1
Unused	Fabric	Fabric	PCIe_1	PCIe_1	x2
Unused	Fabric	Fabric	Fabric	Fabric	Unused

Note: Fabric includes PMA, 8b/10b, 64b/66b, 64b/67b, and PIPE modes.

4.3.2 Inputs/Outputs

PolarFire SoC FPGA I/Os are grouped into pairs to meet the differential I/O standards. Additionally, they are grouped in lanes of 12 buffers with a lane controller for memory interfaces, as shown in the following illustration.

Figure 4-4. I/O Topology



The number of I/O pins varies depending on the device and package size. The persistent I/O feature preserves a state on an I/O without user intervention during programming. The PolarFire FPGA I/O buffers are constructed from the following main sub modules.

- Transmit buffer (PVT compensated)
- Receive buffer
- Termination (Thevenin, Differential, Up, and Down)
- Weak pull mode logic (Up, Down, and Bus-Hold)

Each I/O is configurable and can comply with a large number of I/O standards. The following are two types of user FPGA I/Os in PolarFire SoC FPGAs.

- High-speed I/O (HSIO) optimized for DDR4 memories at speeds up to 1.6 Gbps and a maximum voltage of 1.8 V nominal
- GPIO capable of supporting multiple standards including 3.3 V with an integrated CDR to support SGMII Ethernet applications

The following table summarizes the single-ended I/O support. These are the unterminated standards used in the transceiver lane protocols. Each I/O supports weak pull-up, pull-down, and bus-keeper options. Additionally, each GPIO has a programmable clamp (that is, ON/OFF). For HSIO, the clamp is always ON.

The following table lists the GPIO LVTTTL or LVCMOS receivers that are also designed to support a limited mixed mode of operation to provide greater board I/O design flexibility. For example, if VDDIO is set to 3.3 V, the I/O receivers can operate at the lower voltage of JEDEC standards.

Table 4-3. GPIO Mixed Receiver Mode Operation Capability

VDDIO (V)	LVCMOS33	LVCMOS25	LVCMOS18	LVCMOS15	LVCMOS12
3.3	Yes	Yes	Yes	Not available	Yes
2.5	Yes	Yes	Yes	Yes	Yes
1.8	Yes	Yes	Yes	Yes	Yes
1.5	Yes	Yes	Yes	Yes	Yes
1.2	Yes	Yes	Not available	Yes	Yes

The following table lists the HSIO mixed receiver mode capability.

Table 4-4. HSIO Mixed Receiver Mode Capability

VDDIO (V)	LVCMOS18	LVCMOS15	LVCMOS12
1.8	Yes	Yes	Yes
1.5	Yes	Yes	Yes
1.2	Not available	Yes	Yes

4.3.3 I/O Digital

The PolarFire FPGA I/O digital logic is used to interface between the FPGA fabric and the I/O buffers. It interfaces between the high-speed I/O buffers and lower-speed FPGA fabric. The I/O digital block consists of the following:

- Delay chain for input or output delay
- Registers and control logic for input modes and output modes

The I/O digital registers can be configured for both input and output DDR and shift register modes and combined DDR-shift register modes. It allows gearing up the output data rate and gearing down the input data rate. The PolarFire FPGA I/O digital logic works in conjunction with fast and low-skew clock distributions optimized for DDR applications, special clock dividers, and other support circuits to guarantee clock domain crossings.

4.3.3.1 I/O Digital Features

The following are I/O digital features.

- Programmable input and/or output delay chain
- Data eye monitor for detecting margin-to-clock edges
- Data eye position optimizer
- Up to 10:1 input deserialization
- Up to 10:1 output serialization
- Support for DDR and SDR interfaces
- Receive slip control to facilitate word alignment
- Fast and low-skew lane clocks per 12 I/Os
- Clock recovery for SGMII and similar interfaces (one per 12 I/Os)
- Flash*Freeze support

4.3.3.2 I/O Digital Modes

The following table lists the associated memory interface, I/O data rate, FPGA clock rate, and its applications.

Table 4-5. I/O Digital Modes

Interface	Direction	I/O Data Rate	I/O Clock Rate (MHz)	Gear Ratio	FPGA Clock Rate (MHz)	Applications
DDR4	BiDir	1.6 Gbps	800	8:1	200	Memory interface
DDR3 (L)	BiDir	1.3 Gbps	650	8:1	162.5	Memory interface
LPDDR3	BiDir	1.3 Gbps	650	8:1	162.5	Low-power memory interface
QDRII+	Input/Output	1.1 Gbps	550	8:1	137.5	Low-latency memory interface
RLDRAM3	Input/Output	1.0 Gbps	500	8:1	125	Low-latency memory interface
7:1 LVDS	Input	800 Mbps	400	7:1	114	Flatlink, Cameralink
CDR	Input	1.25 Gbps	625	10:1	125	1000BASE-T, SGMII
MIPI-DPHY	Input/Output	800 Mbps/ 500 Mbps	250	2:1	125	MIPI CSI, DSI
Wide LVDS	Input/Output	1.6 Gbps	800	8, 4, 2:1	250	ADC, DAC

4.4 Non-Volatile FPGA Fabric

The non-volatile FPGA fabric is built on state-of-the-art 28nm low power non-volatile process technology. The PolarFire FPGA fabric is composed of the following building blocks:

- Logic element
- On-chip memory (LSRAM, μ SRAM, sNVM, and μ PROM)
- Math block

The FPGA fabric configuration cells are SEU immune, and are used to configure I/Os and other aspects of the device. Non-volatile FPGAs do not require the configuration process inherent in SRAM FPGAs. Non-volatile FPGAs power up quickly like an ASIC with minimal inrush current, and are ideal for root-of-trust, first-up functionality in any system.

4.4.1 Logic Element

The 4-input LUT can be configured either to implement any 4-input combinatorial function or to implement an arithmetic function where the LUT output is XORed with a carry input to generate the sum output.

The logic element has the following features.

- A fully permutable 4-input LUT optimized for lowest power
- A dedicated carry chain based on a carry look-ahead technique
- A separate flip-flop that can be used independently from the LUT

4.4.2 On-Chip Memory

PolarFire FPGAs integrate four different types of memories that allow the designer to optimize for power, functionality, and area. Two memory types are volatile, and two memory types are non-volatile.

Volatile memories:

- LSRAM
- μ SRAM

The LSRAMs are 20 Kbytes SRAMs with a built-in SECDED and interleaving to prevent multi-bit upsets (MBUs). The μ SRAMs are small distributed 64 x 12 RAMs, well suited for efficient implementation of small buffers, thereby reserving LSRAM usage for the wider and deeper memories.

Non-volatile memories (NVMs):

- μ PROM
- sNVM

The μ PROM, constructed of SEU-immune, FPGA-configuration non-volatile cells, is readable at run time and writable during device programming. It provides users with SEU-immune parameters, constants, IDs, and parametric or initialization data. The sNVM is accessible through system service calls. Data written to the sNVM can be protected by the PUF. The sNVM is readable and writable by the designer's application during runtime and is an ideal storage location for the boot code for soft processors and user keys.

4.4.3 LSRAM

Each LSRAM block consists of 20,480 bits of RAM and includes functionality to support dual-port and two-port modes. There are numerous configurations and features for each block. The Libero SoC PolarFire toolset has an LSRAM configurator that provides automated combining and cascading of several LSRAM blocks into larger memories.

LSRAM features include:

- 428 MHz operation
- True dual-port memory
- Two-port memory (one dedicated write port, one dedicated read port)
- Data widths of $\times 1$, $\times 2$, $\times 5$, $\times 10$, $\times 20$, $\times 40$, and $\times 33$ with SECDED enabled
- Multi-bit-upset mitigation
- Synchronous operation
- Independent port clocks
- Byte enables
- Registered inputs
- Output registers with separate enables and synchronous resets
- Read enables to conserve power while retaining output data
- Power switch to minimize static power when the LSRAM is not used
- Fast zeroization mode

4.4.3.1 Dual-Port Mode

In dual-port mode, the width of both ports is less than 33 and the ports are independent of each other. The write and read operations can occur independently of each other, at any location. On write collisions, while the write operations occur correctly, the read operations can return ambiguous results while the write completes. After completing the write operation, the read data reads the newly written write data correctly.

4.4.3.2 Two-Port Mode

In two-port mode, at least one port has a width of 32 or 40 (or 33 with SECDED). Port A is dedicated for reads and port B is dedicated for writes.

The following illustration shows port widths in various modes.

Figure 4-5. LSRAM Dual- and Two-Port Configurations

	Port A Width						
	x1/x1	x1/x2	x1/x4	x1/x8	x1/x16	W1/R32	N/A
Port B Width	x2/x1	x2/x2	x2/x4	x2/x8	x2/x16	W2/R32	N/A
	x4/x1	x4/x2	x5/x5	x5/x10	x5/x20	W5/R40	N/A
	x8/x1	x8/x2	x10/x5	x10/x5	x10/x20	W10/R40	N/A
	x16/x1	x16/x2	x20/x5	x20/x10	x20/x20	W20/R40	N/A
	W32/R1	W32/R2	W40/R5	W40/R10	W40/R20	W40/R40	N/A
	N/A	N/A	N/A	N/A	N/A	N/A	Wx33/R33
Dual Port							
Two Port							
Two Port SECEDED							

4.4.4 μ SRAM

The μ SRAM is a two-port memory embedded in the FPGA fabric, which is provided for an efficient low-power implementation for small buffers. On write collisions, the write operations occur correctly, while the read operations can return ambiguous results while the write completes. After completing the write operation, the read data reads the newly written write data.

The following are key features of the μ SRAM block:

- 480 MHz operation
- Two-port memory with 64 words of 12 bits
- The write port operates synchronously
- The write port has a fixed width
- The read port operates asynchronously and supports synchronous and pipeline operations with the FPGA fabric flip-flops
- The Libero SoC PolarFire toolset provides automated combining and cascading for larger memories
- Multiple memory blocks can be combined to extend the depth or width
- Provides a state-keeping, low-power suspend mode
- Implemented as an array of latches

4.4.5 μ PROM

The μ PROM is a single monolithic non-volatile memory that provides a PROM-like storage for a variety of purposes, including initialization data for other memories, user calibration data, and so on. The memory cells are constructed from the FPGA configuration cells and are updated when the device is programmed.

The following are key features of the μ PROM:

- 10 ns read access time
- Programmed with the FPGA bitstream
- Asynchronous or synchronous read access mode from the FPGA fabric

4.4.6 sNVM

Each PolarFire FPGA has 56 Kbytes of sNVM. The sNVM is organized into 221 pages of 236 bytes or 252 bytes, depending on whether the data is stored as plain text or encrypted/authenticated data. It is accessible to users through system services calls to the PolarFire FPGA system controller. Pages within the sNVM can be marked as ROM during bitstream programming. The sNVM content can be used to initialize LSRAM and μ SRAMs with secure

data. The sNVM is only accessible through system service calls. Data written to the sNVM can be protected by the PUF.

4.4.7 Math Block


The fundamental building block in any digital signal processing algorithm is the multiply-accumulate (MACC) operation. PolarFire FPGAs implement a custom 18 x 18 MACC block for an efficient, low-power implementation of complex DSP algorithms such as finite impulse response (FIR) filters, infinite impulse response (IIR) filters, and fast Fourier transform (FFT) for filtering and image processing applications. An optional 16-word coefficient ROM can be constructed from logic elements located near the math block.

The following are key features of the math block functionality:

- 545 MHz operation
- 18 × 18 two's complement multiplier accumulator with an output width of 48 bits
- Power-saving pre-adder to optimize linear phase FIR filter applications and reduce the math block usage
- Optional pipelining and dedicated buses for cascading
- Dot-product mode for complex multiplies

The following illustration shows the functional blocks of the math block.

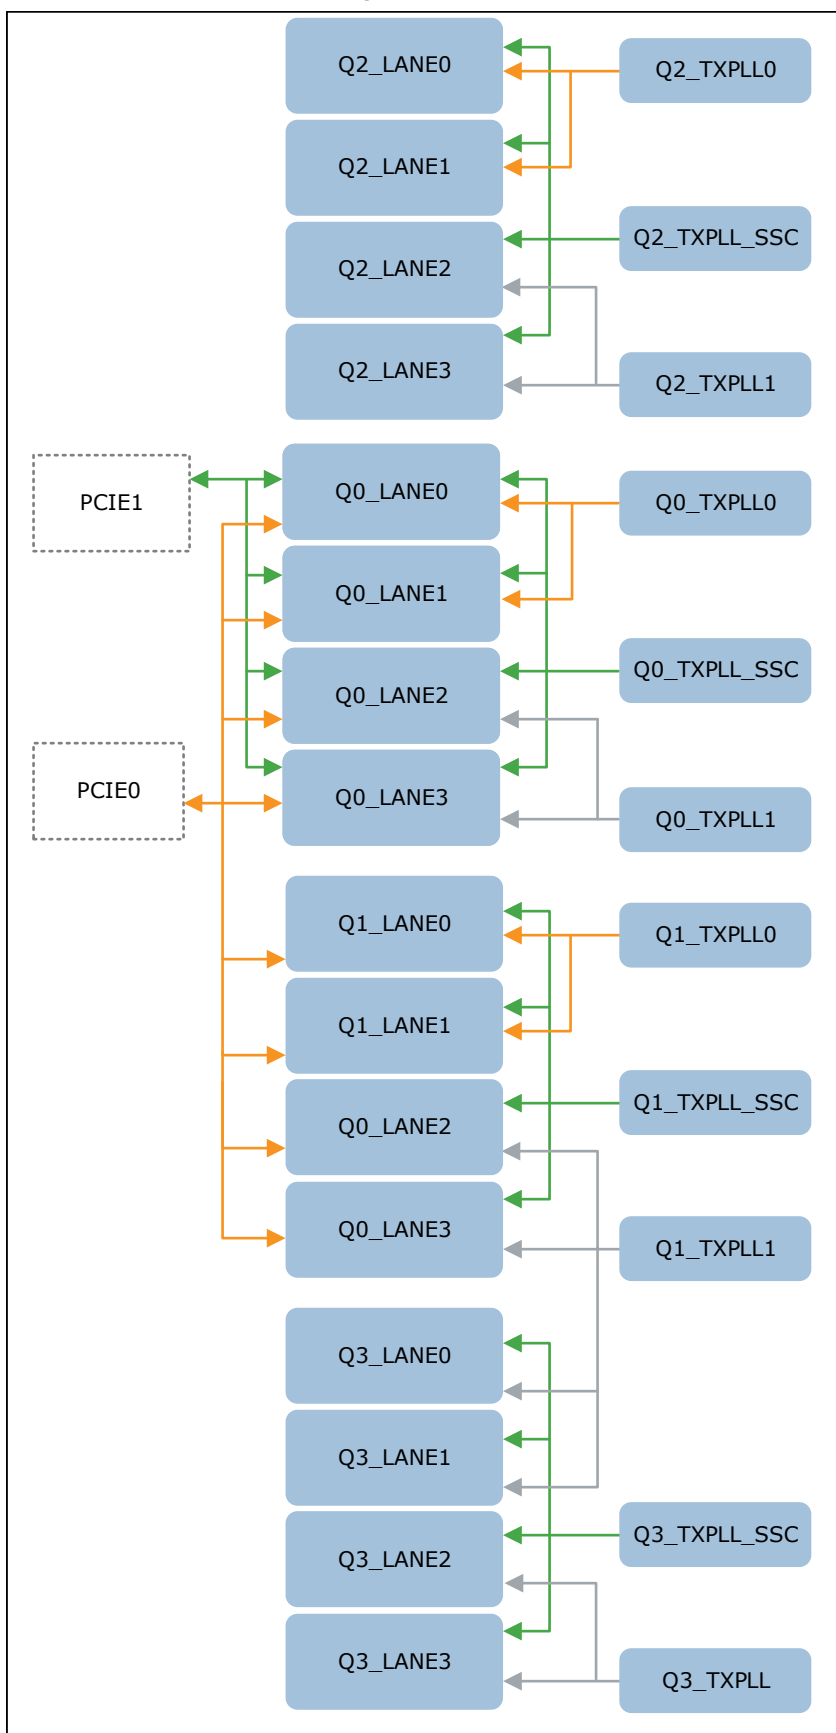
Figure 4-6. Math Block



4.5 PCI Express

Each PolarFire FPGA integrates two low-power, built-in PCIe Gen2 controllers, allowing seamless and easy connectivity to one or more host processors. The two PCIe controllers are shared across two quads, as shown in the following illustration. All PLLs are jitter attenuation-capable, while the SSC label indicates spread spectrum clock (SSC) capability.

Figure 4-7. PCI Express Hard Macro Lane Sharing



4.5.1 PCI Express Features

The following are PCIe features.

- ×1, ×2, and ×4 lane support
- Suitable for root port, native endpoint
- PCI Express base specification revision 2.0 and 1.1 compliant
- AXI4 master and slave interfaces to the FPGA fabric
- Single function capability
- Advanced error reporting (AER) support
- Integrated clock domain crossing (CDC) to support user-selected AXI4 frequency
- Lane reversal support
- Legacy PCI power management support
- Native active state power management L0s and L1 state support
- Power management event (PME message)
- MSI and legacy INT message support
- Latency tolerance reporting (LTR)
- L1 PM sub-states with CLKREQ
- Address translation tables between the PCIe and AXI4 domains

4.5.2 PCI Express DMA Engines

Each PCIe controller supports the following built-in DMA modes, enabling low-power and efficient data transfer into the FPGA fabric.

- Two DMA channels
- Eight outstanding read and write requests
- Completion reordering support
- Flexible scatter-gather DMA modes, including dynamic DMA control per descriptor
- Optional DMA engine reporting to the descriptor to ease software management
- Fetching of up to three descriptors to optimize throughput

5. System Controller

The PolarFire SoC FPGA system controller is based on the industry-standard ARM Cortex-M3 and is only used for FPGA power-up, secure DPA-safe FPGA programming, and executing and responding to system services. All internal memories are SECDED-protected with background scrubbing capabilities to remove single-bit errors.

5.1 System Services

System services provide the user with information about the state of the FPGA and allow the user to request the system controller to perform predefined functions using a standard Application Programming Interface (API).

The system services are listed as follows.

Design Services

- Initialize fabric RAM
- Bitstream authentication
- IAP image authentication
- Device services
- Serial number
- JTAG user code
- Design version number
- Device certificate

Data Services

- sNVM read/write
- PUF emulation service
- Nonce service
- FPGA fabric services
- Digest check
- In-application programming

5.2 Programming

PolarFire SoC FPGAs have multiple programming modes designed to enable various use models. All bitstreams are always encrypted and DPA safe. Each PolarFire SoC FPGA can be programmed using a dedicated SPI peripheral and JTAG port. All PolarFire SoC FPGAs are typically reprogrammed in less than 60 seconds. For device specific programming timings, see the PolarFire SoC FPGA Datasheet.

The following programming modes are supported:

Slave Programming

- JTAG
- Slave SPI—an external SPI master programs the FPGA

SPI Master Programming—In-Application Programming (IAP)

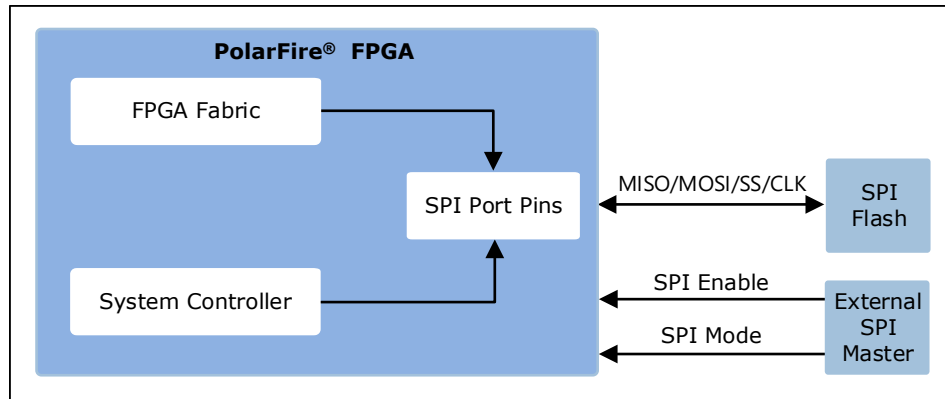
- Auto update feature—the system controller on power-up checks for a new bitstream in an external SPI flash and programs the FPGA.
- Auto programming feature—on a blank device, the system controller on power-up checks for a bitstream in an external SPI flash and programs the FPGA.
- Programming recovery feature—if remote programming fails due to a power interruption, the system controller reprograms the FPGA on the next power-up cycle from a golden bitstream (located in an external SPI flash).

5.2.1 Dedicated SPI Programming Port

To facilitate the use of various programming modes, PolarFire SoC FPGAs share dedicated SPI port pins between the system controller and user logic embedded in the FPGA. User logic must instantiate the User SPI macro to gain access to the pins from their design. The SPI port pins can be used as a master or slave programming port based on the signal level on the dedicated SPI mode pin. The dedicated SPI Enable pin also allows an external SPI master to program the on-board SPI flash without an external MUX by tri-stating the SPI MOSI/MISO/SS/CLK pins on the PolarFire FPGA.

The following illustration shows the SPI port facilitating the use of various programming modes.

Figure 5-1. SPI Programming Port



6. Low Power

PolarFire SoC FPGAs offer a variety of techniques and capabilities to lower the total application power. Users can take advantage of these features to lower both capital and operational expenditures with smaller or no heat sinks, smaller or fewer fans, lower cooling costs, and so on. Additionally, the lower total power advantage can also allow the user to pack more compute operations into an existing thermal budget.

6.1 Non-Volatile Technology

Using a non-volatile complementary metal–oxide semiconductor (CMOS) technology for the FPGA configuration cells offers several power advantages over SRAM FPGA technology.


- A non-volatile switch has lower power than a SRAM switch, leading to lower static power consumption
- No SRAM configuration in-rush currents
- An external configuration component is not necessary

6.2 Low-Power Transceiver Lane

PolarFire SOC FPGAs' low-power capability is also extended to the industry's most power efficient transceiver lane, enabling 10GBASE-KR applications at less than 100 mW of power per lane. The transceiver lane has comprehensive power-down controls to optimize power consumption, including programmable amplitude and edge rate control.

The following illustration shows the connection between transceiver power and data rate.

Figure 6-1. Transceiver Power versus Data Rate



7. Reliability

Microchip continues to offer the industry's most reliable FPGAs for your mission and safety critical applications.

7.1 FPGA Fabric

PolarFire SoC FPGA configuration cells are inherently immune to SEUs caused by neutrons. Contrary to popular belief, shielding does not prevent a neutron from passing through an electronic system or electronic device. As semiconductor device geometry shrinks to smaller lithography, the problem of MBUs starts appearing. SRAM FPGA scrubbing techniques might be inadequate in these circumstances and while scrubbing may help, an important point is that scrubbing detects an error after the fact. The error has already occurred and propagated throughout the system. The configuration of the PolarFire SoC FPGA fabric provides worry-free operation against random events caused by SEUs.

7.2 LSRAM

LSRAMs have built-in SECEDED capability on a 32-bit word boundary. Seven additional bits are used for error correction. Two flags are provided to the user to indicate SECEDED. Mitigation against multi-bit upsets is provided by keeping all cells in a word separated by a minimum distance. Applications that require scrubbing need to be accomplished with user logic. The error correction logic can be turned ON and OFF by the user to enable easy validation of the error correction operation.

7.3 μ SRAM

The 64×12 μ SRAMs are constructed from latches and are not as sensitive to SEUs as SRAMs are.

7.4 Digests

Digests verify the integrity of the programmed non-volatile data. Digests are a cryptographic hash of various data areas. Any digest that reports back an error raises the digest tamper flag.

The following are digestible non-volatile areas:

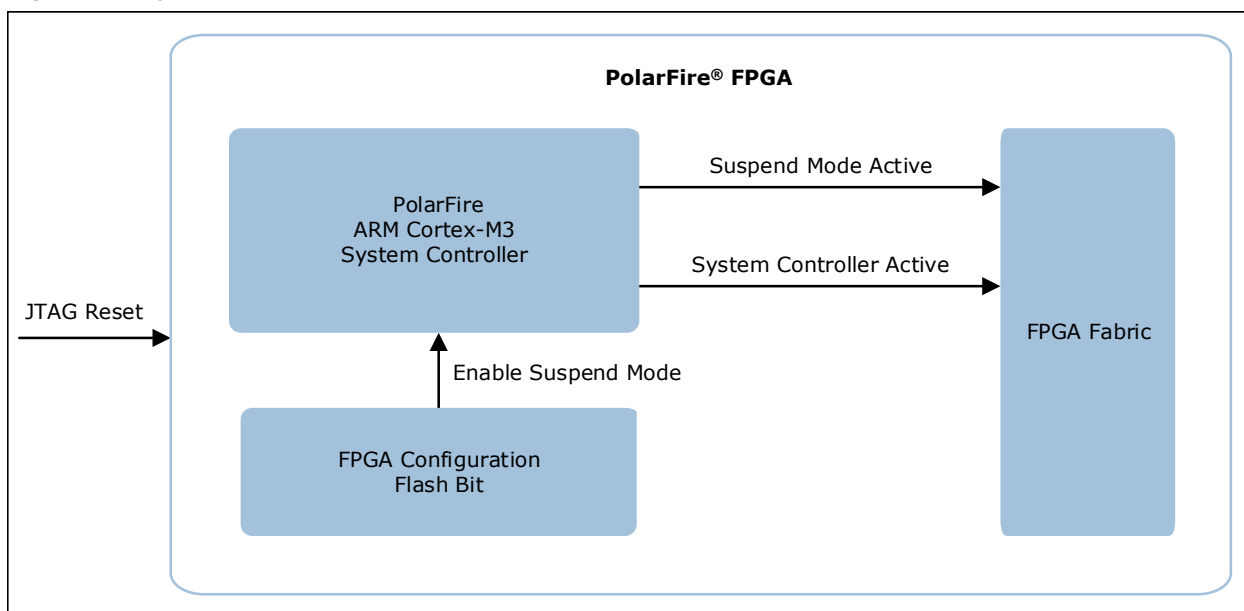
- The FPGA fabric and consequently the μ PROM
- sNVM marked as ROM
- User key 1
- User key 2
- Factory parametric and key storage
- 128 Kbytes eNVM block

7.5 System Controller Suspend Mode

For safety critical applications, PolarFire SoC FPGAs allow the user to place the Cortex-M3-based system controller in a reset state after the FPGA has powered up. By programming an SEU configuration non-volatile bit, the Cortex-M3 is placed in reset by a TMRed SEU immune reset latch after FPGA power-up. User logic can monitor if the suspend mode command is active and if the system controller cannot fetch instructions while in the reset state. The FPGA can be re-programmed after disabling the suspend mode by asserting the appropriate JTAG signals. The JTAG TRSTB signal must be asserted low for suspend mode to remain active.

The following illustration shows how to activate and deactivate suspend mode.

Figure 7-1. System Controller Suspend Mode



8. Security

Today's demanding applications not only have to meet the functional requirements, but also to meet them in a secured way. Security starts during silicon manufacturing and continues through system deployment and operations.

8.1 Design Security

Protecting your design starts with wafer manufacturing and continues through the deployment of the end product. The following are key features that provide state-of-the-art supply chain assurance and IP protection benefits:

- Secure supply chain management through the use of hardware security modules (HSMs) during wafer test and packaging
- Supply chain assurance through the use of a 768-byte digitally signed x.509 FPGA certificate embedded in every FPGA
- AES256-encrypted CRI DPA countermeasures patent protected, bitstream, and key management protocols
- Built-in tamper detectors: voltage monitors, temperature monitor, clock glitch detectors, voltage glitch detectors, protective meshes, and bus scrambling
- Data integrity through built-in cryptographic digest capabilities
- Zeroization capabilities for all on-chip memories and the FPGA fabric
- Integrated PUF for key storage
- 56 Kbytes of PUF protected sNVM
 - Secure reprogrammable keys using non-volatile memory

8.1.1 Tamper Detectors

PolarFire SoC FPGAs integrate numerous on-chip tamper detectors, enabling users to monitor the environment and the operating parameters of the design. The user can respond to the events that are determined to be out-of-scope for proper operation. Tamper flags indicate that a tamper event has occurred and are available as signals to the FPGA fabric for users to process and respond. The following is a partial list of tamper detectors:

- Clock glitch detectors
- Clock frequency detectors
- Voltage monitor detectors
- Temperature sensor
- JTAG active detector
- Mesh active detector

8.1.2 Tamper Responses

After processing a detected event, the user can perform one of the following actions.

- Disable I/Os—configurable on a per I/O basis
- Security lockdown
- Reset
- Zeroize

9. PolarFire SoC Device Offerings

PolarFire SoC FPGAs offer low-power transceiver devices and various device offerings with transceivers, such as design security, and low-power data security. All PolarFire SoC devices are integrated with multi-protocol industry-leading low-power transceivers. Low power (L) devices provide up to 35 percent lower static power.

The following table lists the extended commercial and industrial PolarFire SoC offerings using the MPFS250T as an example. The MPFS025T, MPFS095T, MPFS160T, and MPFS460T device densities have identical offerings. Temperatures listed are junction temperatures.

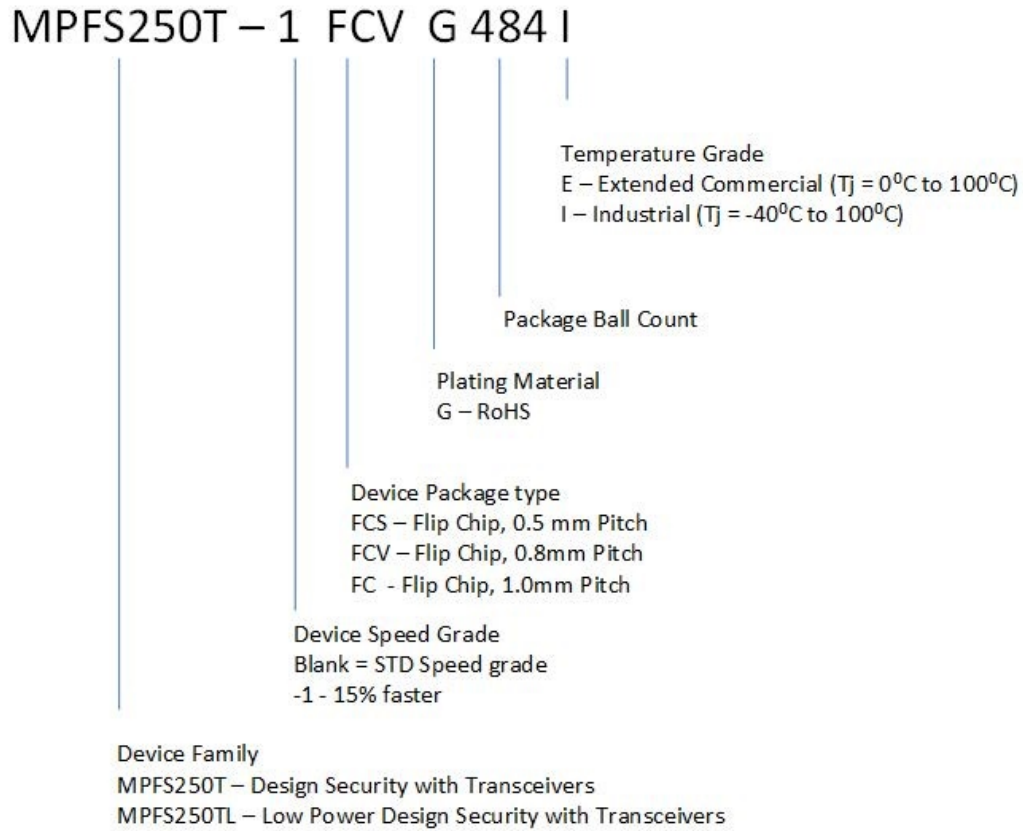
Table 9-1. PolarFire SoC Offerings

Device Options	Extended Commercial Temperature (E) 0 °C–100 °C	Industrial Temperature (I) –40 °C–100 °C	STD Speed Grade	–1 Speed Grade	Transceivers (T)	Lower Static Power (L)
MPFS250T	Yes	Yes	Yes	Yes	Yes	—
MPFS250TL	Yes	Yes	Yes	—	Yes	Yes

10. Ordering Information

PolarFire SoCs are offered with multiple speed grades, temperatures, and package combinations. All temperatures are specified as junction temperatures. The following illustration shows the ordering information.

Figure 10-1. Ordering Information



11. Revision History

Revision	Date	Description
A	09/2020	In Revision A, the document was updated to Microchip template.
1.0	12/2019	This is the initial release of this document.

The Microchip Website

Microchip provides online support via our website at www.microchip.com/. This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

Product Change Notification Service

Microchip's product change notification service helps keep customers current on Microchip products. Subscribers will receive email notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, go to www.microchip.com/pcn and follow the registration instructions.

Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: www.microchip.com/support

Microchip Devices Code Protection Feature

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods being used in attempts to breach the code protection features of the Microchip devices. We believe that these methods require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Attempts to breach these code protection features, most likely, cannot be accomplished without violating Microchip's intellectual property rights.
- Microchip is willing to work with any customer who is concerned about the integrity of its code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is "unbreakable." Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Legal Notice

Information contained in this publication is provided for the sole purpose of designing with and using Microchip products. Information regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications.

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL OR CONSEQUENTIAL LOSS, DAMAGE, COST OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AnyRate, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Klear, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PackeTime, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TempTrackr, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, FlashTec, Hyper Speed Control, HyperLight Load, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, Vite, WinPath, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, EtherGREEN, In-Circuit Serial Programming, ICSP, INICnet, Inter-Chip Connectivity, JitterBlocker, KlearNet, KlearNet logo, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, SAM-ICE, Serial Quad I/O, SMART-I.S., SQI, SuperSwitcher, SuperSwitcher II, Total Endurance, TSHARC, USBCheck, VariSense, ViewSpan, WiperLock, Wireless DNA, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2020, Microchip Technology Incorporated, Printed in the U.S.A., All Rights Reserved.

ISBN: 978-1-5224-6765-6

Quality Management System

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
Corporate Office 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Fax: 480-792-7277 Technical Support: www.microchip.com/support Web Address: www.microchip.com	Australia - Sydney Tel: 61-2-9868-6733 China - Beijing Tel: 86-10-8569-7000 China - Chengdu Tel: 86-28-8665-5511 China - Chongqing Tel: 86-23-8980-9588 China - Dongguan Tel: 86-769-8702-9880 China - Guangzhou Tel: 86-20-8755-8029 China - Hangzhou Tel: 86-571-8792-8115 China - Hong Kong SAR Tel: 852-2943-5100 China - Nanjing Tel: 86-25-8473-2460 China - Qingdao Tel: 86-532-8502-7355 China - Shanghai Tel: 86-21-3326-8000 China - Shenyang Tel: 86-24-2334-2829 China - Shenzhen Tel: 86-755-8864-2200 China - Suzhou Tel: 86-186-6233-1526 China - Wuhan Tel: 86-27-5980-5300 China - Xian Tel: 86-29-8833-7252 China - Xiamen Tel: 86-592-2388138 China - Zhuhai Tel: 86-756-3210040	India - Bangalore Tel: 91-80-3090-4444 India - New Delhi Tel: 91-11-4160-8631 India - Pune Tel: 91-20-4121-0141 Japan - Osaka Tel: 81-6-6152-7160 Japan - Tokyo Tel: 81-3-6880-3770 Korea - Daegu Tel: 82-53-744-4301 Korea - Seoul Tel: 82-2-554-7200 Malaysia - Kuala Lumpur Tel: 60-3-7651-7906 Malaysia - Penang Tel: 60-4-227-8870 Philippines - Manila Tel: 63-2-634-9065 Singapore Tel: 65-6334-8870 Taiwan - Hsin Chu Tel: 886-3-577-8366 Taiwan - Kaohsiung Tel: 886-7-213-7830 Taiwan - Taipei Tel: 886-2-2508-8600 Thailand - Bangkok Tel: 66-2-694-1351 Vietnam - Ho Chi Minh Tel: 84-28-5448-2100	Austria - Wels Tel: 43-7242-2244-39 Fax: 43-7242-2244-393 Denmark - Copenhagen Tel: 45-4485-5910 Fax: 45-4485-2829 Finland - Espoo Tel: 358-9-4520-820 France - Paris Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79 Germany - Garching Tel: 49-8931-9700 Germany - Haan Tel: 49-2129-3766400 Germany - Heilbronn Tel: 49-7131-72400 Germany - Karlsruhe Tel: 49-721-625370 Germany - Munich Tel: 49-89-627-144-0 Fax: 49-89-627-144-44 Germany - Rosenheim Tel: 49-8031-354-560 Israel - Ra'anana Tel: 972-9-744-7705 Italy - Milan Tel: 39-0331-742611 Fax: 39-0331-466781 Italy - Padova Tel: 39-049-7625286 Netherlands - Drunen Tel: 31-416-690399 Fax: 31-416-690340 Norway - Trondheim Tel: 47-72884388 Poland - Warsaw Tel: 48-22-3325737 Romania - Bucharest Tel: 40-21-407-87-50 Spain - Madrid Tel: 34-91-708-08-90 Fax: 34-91-708-08-91 Sweden - Gothenberg Tel: 46-31-704-60-40 Sweden - Stockholm Tel: 46-8-5090-4654 UK - Wokingham Tel: 44-118-921-5800 Fax: 44-118-921-5820

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

Microchip:

[MPFS250T-FCG1152EES](#) [MPFS250T-1FCG1152EES](#) [MPFS250T-1FCSG536EES](#) [MPFS250T-1FCVG484EES](#)
[MPFS250T-1FCVG784EES](#) [MPFS250T-FCSG536EES](#) [MPFS250T-FCVG484EES](#) [MPFS250T-FCVG784EES](#)