



BLAZAR BE2-BURST Accelerator Engine

Intelligent In-Memory Computing

576Mb Memory



PRODUCT BRIEF: MSR622

Acceleration Engines Give Software and Hardware System Architects Acceleration Options not Previously Available

Bandwidth Engine (BE) Overview

The **BLAZAR Family of Accelerator Engines** are high capacity, high-speed memories that support high bandwidth, fast random memory access rates and include optional *embedded In-Memory Functions (IMF)* that solve critical memory access challenges for memory bottlenecked applications like network search, statistics, buffering, security, firewall, 8k video, anomaly detect, genomics, ML random forest of trees, graph walking, traffic monitoring, AI and IoT.

All Accelerator Engines have two ways to be uses in a system:

1. As a standard parallel like QDR, SyncSRAM or RLDram
2. As a high density, high bandwidth memory with Optional In-Memory Acceleration Functions

Both modes are independent. You do not have to use the In-Memory functions if they are not useful in your application. In that case, you would use the Accelerator Engine like any other memory.

Base Features: BE2-BURST

- 576 Mb of tRC of 3.2 ns memory
 - Replaces 4 QDR type memories
- In-Memory BURST functions
- RTL Memory Controller(s) available

Application Focus

- Slower speed applications needing high capacity
- SRAM with high capacity and high speed
- High bandwidth data access application where low latency and movement of data is critical
- FPGA Acceleration for Xilinx and Intel

Key Features: Memory

- 576 Mb SRAM (8M x 72b)
 - User defined WORD width
 - Typical 8x, 16x, 32x, 36x, ... 72x
- High bandwidth, low pin count serial interface
 - Highly efficient reliable transport command and data protocol optimized for 90% efficiency
 - Eases board layout and signal integrity, minimal trace length matching required, operates over connectors
 - Reduction of I/O pins from 5x to 20x depending on equivalent memory density and type
- High access rate SRAM class memory
 - Up to 3.3 Billion transactions/sec
 - 3.2 ns tRC
- Highest single chip bandwidth – up to 320 Gb/s throughput (160 full duplex)

Key Features: In-Memory Functions

- The acceleration function is optional and does not impact the device used as a memory only
- BURST In-Memory function
 - For sequential Read or Write functions for data movement
 - Burst length: 1, 2, 4 or 8 words
 - Can double or triple QDR bandwidth

MoSys Accelerator Engine Elements of BE2-BURST

MoSys Engines have a Unique Memory Architecture that can replace SyncRAM/RLDRAM memories and Embeds In Memory Functions (IMF) that execute many times faster. A single function replaces many traditional memory accesses.



Bandwidth Engine Device Key Features

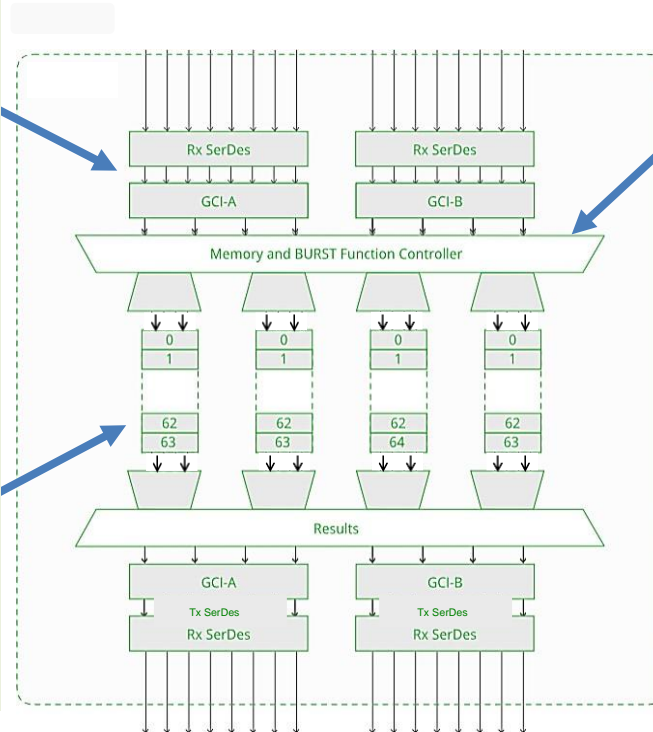
BE2-BURST

High-Speed Serial I/O

- GCI serial I/O versions of 10 and 12.5 Gbps for high bandwidth (up to 320 Gbps)
- Device can operate with a minimum of 4 lanes.
- Has two, full duplex 8 lane ports that operate independently
- Reduces number of signal pins over traditional memories, increases signal integrity allowing longer board traces to ease board signal routing
- Operates across connector

Main Memory

- 576Mb (BE3 has 1Gb)
 - 4 partitions/64 banks
 - 8 READ & 8 WRITE ports
- 3.2 ns tRC
- Allows parallel partition and bank execution



Memory/Function Controller

- Directs all read/write function execution to selected bank of memory
- Manages the sequence of In- Memory functions
 - BURST – Sequential read or writes
 - Up to 8 read and/or writes
- Controls simultaneous memory access to partitions and banks

Common Key Features for Bandwidth Engines (BE2 & BE3)

- Tx SerDes High capacity, high-speed memory on a single device
 - BE2 576Mb
 - BE3 1Gb
- High speed tRC access
 - BE2 with 3.2ns
 - BE3 with 2.7ns
- Achieves the highest bandwidth possible
 - Use of serial interface with high efficiency GCI protocol
 - BE2 with 3.3B access per sec
 - BE3 with 6.5B accesses per sec
- Reduces number of interface pins compared to other memories
 - Minimum use of 4 lanes on one port or 16 pins
 - Typical system uses 8 lanes or 32 pins
 - Highest bandwidth uses 16 lanes or 64 pins
- Eliminates external components required for signal integrity
 - On device Auto-Adaption to eliminate external signal conditioning board components
 - Signals will work over a backplane
- RTL makes the interface like a parallel QDR
 - MoSys supplies an RTL memory controller that handles the memory and serial interface
 - Serial device interface is transparent to user
 - Provides a parallel QDR like Read/Write RTL interface
- Make interface easily adapted to an AXI or Avalon bus
 - Minimal RTL logic required

Benefits of BE2 vs QDR


Summary of Benefits

- Capacity ... 576Mb memory ... Replaces 4 QDR/SyncSRAM devices
- Cost... One BE-2 with 4x the capacity is approximately the price of two QDR memories
- Pins ... Typical application uses only 8 signals (32 pins) with signal auto-adaptation

More High-Speed memory generally allows acceleration options for software and hardware architects/designers

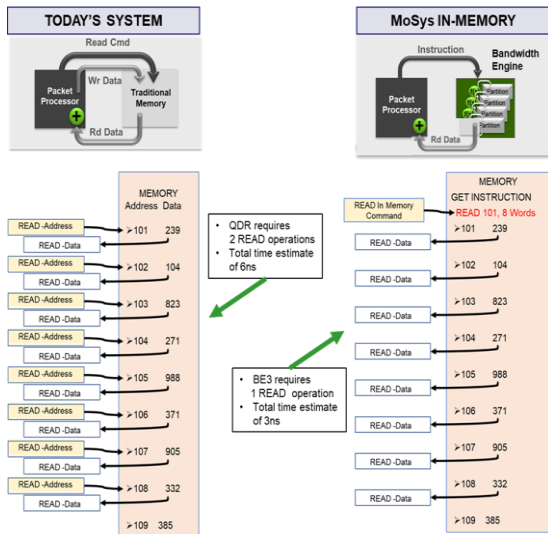
Overview Comparison BE vs. QDR

- Memory size
 - BE2 with 576Mb equivalent to 4 QDRs with 144Mb per device
- Device PCB board space saving
 - 1 BE2 device vs 4 QDR devices
- Signal pins reductions
 - 4 QDR...576Mb requires 500-720 pins
 - 1 BE2 ...576Mb...BE2 typical system uses 8 lanes or 32 pins
 - All BE devices have Auto-Adaptation which handles on-board signal tuning, eliminating the need for any external components to insure a clean, reliable signals
- Cost
 - One BE-2 with 4x the memory capacity is approximately the price of two QDR memories
- Application Benefits
 - Larger buffers, High Bandwidth
 - Allows real time operations and analysis at line rate
 - Eliminates need for complex parallel operations using RLDRAM, HBM, or slow DRAM



In-Memory Functions

Example of BURST Function (BE2/BE3/PHE)



When there is a need to move data at a high bandwidth, the In-Memory commands can save a tremendous amount of time in one tRC cycle

In the BE2,

- 8 Read
- 8 Writes
- 8 Reads + 8 Writes

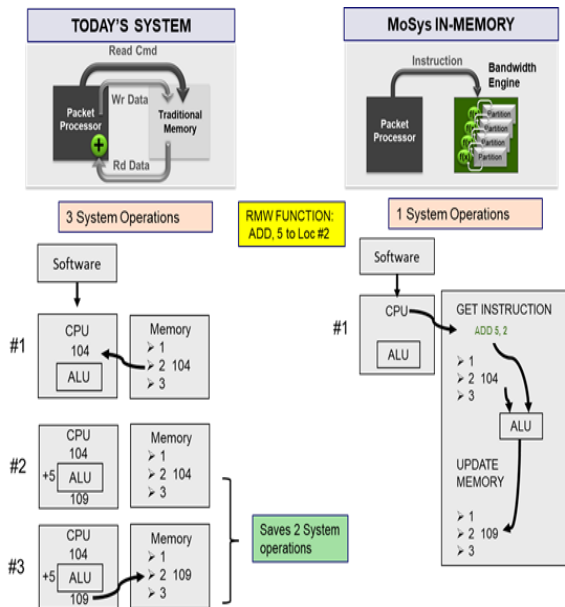
In the BE3

- 16 Reads
- 16 Writes
- 16 Reads + 16 Writes

Example In-Memory BURST time saving

- tRC 3ns
- QDR 144b read
- QPR 576b read
- QPR 4x of a QDR
- There are more than 12 different commands

Example of RMW Function (BE2/BE3/PHE)



Focused on DATA COMPUTING AND DECISION where there is need for memory location modification involving RMW in applications such as metering, as well a single or dual counter update for statistics.

There are over 27 operations available such as add, subtract, compare, increment, etc.

Example In-Memory RMW time saving

Add a Number to a Location (RMW)

QDR Traditional Memory System

- 3 operations Time Analysis
- Total Time = 6ns + FPGA ADD TIME

MoSys In Memory Function

- 1 operations Time Analysis
- Total Time = 3ns

Compatibility Quazar - Blazar Family of Accelerator Engines

- MSP220 (QPR4) pin compatible MSR622/MSR820 (BE2)
- MSP230 (QPR8) pin compatible MSR630/MSR830 (BE3)

Simplifying User Interface to BE2/3 w/MoSys RTL Controller


Interface Ports



RTL supplied by MoSys makes the serial interface transparent by converting it to an RTL Parallel QDR like interface

- Device has two, 8 lane independent ports
- Typical system uses one port (shown), 8 lanes or 32 pins
- Can use as few as 4 lanes on one port or 16 pins
- High bandwidth systems use both ports, 16 lanes, 64 pins
- Independent ports can operate as a dual port with simultaneous access between two FPGAs

Serial Interface




- Benefits of a serial interface allows high bandwidth over very few pins
- Key to bandwidth is the MoSys GCI interface that is transparent to the user with the MoSys-supplied FPGA RTL memory controller

MoSys FPGA Parallel RTL Interface

MoSys-Supplied RTL Controller Simplifies User Interface with the BE

MoSys-supplied FPGA RTL Memory Controller interfaces with the MoSys Bandwidth Engine. This controller is between the User Application RTL logic and the BE device.





MoSys-Supplied RTL Controller Simplifies User Interface with the BE

- It handles all the logic for the Serial GigaChip Interface (GCI) in the FPGA
- Eliminates the user having to design a serial interface by making it transparent and providing a QDR parallel like interface
- Memory word width is user-definable in RTL
 - Typical word widths are 8, 16, 32, 36, 64 ...
- While the memory on the BE2 is organized as 8Mx72b and the BE3 is 16Mx72b, the address conversion mapping from the selected WORD width to the BE memory is handled by the RTL
 - Address translation to BE memory organization is transparent to the application
- All memory addressing and commands are presented to the QDR-like parallel interface.
- If the optional In-Memory Functions are use, the RTL controller will manage their execution

The signal interface at the User Application is a simple SRAM memory address, data, control structure. This simple interface shields the users from the BE commands, serial interface and the scheduling logic and memory partition timing.

High Speed GCI Serial Interface



SIGNAL NAME	WIDTH	DIR	DESCRIPTION
Read Interface			
rd_p	1	In	Assertion of this signal indicates that this is a read transaction.
rd_addr_p	32	In	Read address. Please refer to the Address section of this specification to see the detail of this address field.
rd_partsel_p	1	In	Indicates the BE-2 partition that this read command will be operated upon: 0 = Partition 0 for GCI port A, Partition 1 for GCI port B 1 = Partition 2 for GCI port A, Partition 3 for GCI port B
rd_data_p0	*	Out	Returned data from BE-2 memory. This data is qualified by the "rd_datav_p0" signal
rd_data_p1	*	Out	Returned data from BE-2 memory. This data is qualified by the "rd_datav_p1" signal. Note that rd_data_p1 will only have valid data if rd_data_p0 is valid as well. rd
rd_datav_p0	1	Out	The Memory Controller asserts this signal to indicate the current data in the "rd_data_p0" bus is valid
rd_datav_p1	1	Out	The Memory Controller asserts this signal to indicate the current data in the "rd_data_p1" bus is valid. Note that rd_data_p1 will only have valid data if rd_data_p0 is valid as well
rd_wait_rq_p	1	Out	The Memory controller asserts "rd_wait_rq_p" to indicate that it cannot accept the current read request from user. The User Application should hold all the request signals (rd_p, rd_addr_p ...) until the de-assertion of this signal.

SIGNAL NAME	WIDTH	DIR	DESCRIPTION
Write Interface			
wr_p	1	In	Assertion of this signal indicates that this is a write transaction.
wr_addr_p	32	In	Write address of the memory for this transaction. Please refer to the Address section of this specification to see the detail of this address field.
wr_partsel_p	1	In	Indicates the BE-2 partition that this write command will be operated upon: 0=Partition 0 for GCI port A, Partition 1 for GCI port B 1=Partition 2 for GCI port A, Partition 3 for GCI port B
wr_data_p	*	In	Write data from the User Application logic.
wr_wait_rq_p	1	Out	The Memory controller asserts "wr_wait_rq_p" to indicate that it cannot accept the current write request. The User Application should hold all the request signals (wr_p, wr_addr_p ...) until the de-assertion of this signal.

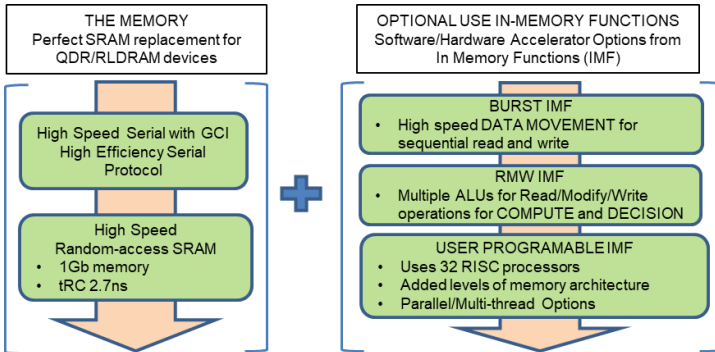
Accelerator Engine Family Overview

Software Defined - Hardware Accelerated

In-Memory	Part Number	Description	Package	Interface				Memory			Access Rate	In-Memory Functions				
			Pkg Size	Lanes	Rate per Lane Gb/s				BW MAX.	tRC	Size	Billion Transactions per second	R/W	BURST for Data Movement	RMW / ALU for Compute and Decision	Custom & User Functions with 32 RISC Cores
			mm	Tx/Rx	10.3	12.5	15.6	25	Gb	ns	Gb					
QPR4	MSQ220	QPR4 (Quad Partition Rate) 0.5 Gb	FCBGA 19X19	16	✓	✓			320	3.2	0.5	2.5	✓			
QPR8	MSQ230	QPR8 (Quad Partition Rate) 1Gb	FCBGA 27X27	16			✓	✓	640		1	5	✓			
BURST	MSR622	Bandwidth Engine 2 Burst Serial 0.5Gb High Access Memory	FCBGA 19x19	16	✓	✓			320	3.2	0.5	3.3	✓	✓		
	MSR630	Bandwidth Engine 3 Burst Serial 1Gb High Access Memory	FCBGA 27x27	16		✓	✓	✓	640	2.7	1	6.5	✓	✓		
RMW	MSR820	Bandwidth Engine 2 RMW Serial 0.5Gb High Access Memory with ALU for RMW functions	FCBGA 19x19	16	✓	✓			320	3.2	0.5	3.3	✓	✓	✓	
	MSR830	Bandwidth Engine 3 RMW Serial 1Gb High Access Memory with ALU for RMW functions	FCBGA 27x27	16		✓	✓	✓	640	2.7	1	6.5	✓	✓	✓	
Program	MSPS30	Programmable HyperSpeed Engine Serial Interface, 1Gb Memory, 32 RISC Processor cores for custom algorithms, compute, functions	FCBGA 27x27	16		✓	✓	✓	717	2.7	1	24 Internal	✓	✓	✓	✓
RTL	RTL-AE	RTL Memory Controller for Bandwidth Engine and Programmable HyperSpeed Engine. Manages memory and the serial interface signals. Presents a QDR like parallel RTL interface to the user.	FPGA RTL Code		✓	✓	✓	✓			576Mb & 1Gb	6.5	✓	✓	✓	

BW MAX. = Aggregate of all SerDes laneS at the highest serial interface speed

In-Memory Acceleration Functions



Optional Function Will Not Impact the Device When Used as Memory Only

BURST In-Memory Function


- For sequential Read or Write functions for **DATA MOVEMENT**
- Burst length: 1, 2, 4 or 8 words
- Can double or triple QDR bandwidth
- Simultaneous execution of read and writes

RMW In-Memory Function


- RMW are Read/Modify/Write functions
- Includes many functions for compute and decision
- Examples: ADD, SUB, Compare, INC plus 15 other functions
- Increases execution speed and bandwidth

Flexible Configuration Uses


CONTACT MOSYS TO LEARN ABOUT THESE ADVANCED FEATURES



DUAL PORT



PIPELINE



HIGH BANDWIDTH

www.mosys.com

<https://mosys.com/products/blazar-family-of-accelerator-engines/>

