Accelerating Test and Measurement with Intelligent Memories

Topics

- **Introduction to MoSys**
  - Memory Tradeoffs vs Performance
- **Test and Measurement**
  - Definitions
  - Needs
  - Block Diag of generalized Test and Measurement systems
- **Accelerator Engines Integrated Circuits**
  - Quazar Family
  - Blazar Family
  - Accelerator Engine Architecture
    - Quad Partition Rate Memory
    - Intelligent In-Memory Function Bandwidth Accelerators
  - Application Examples using Accelerator Engines
- **Programmable HyperSpeed Engine (PHE)(Brief Introduction)**

Accelerator Engines & Linespeed Product Lines

QUAZAR FAMILY of Integrate Circuits

Memory

Performance/Cost Target Memories

Quad Partition Rate Engine (QPR)

QPR4 576Mb memory
QPR8 1Gb memory

Low Cost QDR alternative
- Replace 4 QDRs
- **For less than $200 in volume**

Higher bandwidth
- Up to 240Gb/s

Lower power

Pin compatible with BLAZAR family

BLAZAR FAMILY of Integrate Circuits

Accelerator Engines (AE)

Memories with Embedded Acceleration Engines

Bandwidth Engine (BE)

BE2 576Mb memory
BE3 1Gb memory

QPR **PLUS** Acceleration
In-Memory Acceleration Functions
Two options
- Burst (12+ functions)
- RMW with ALU (17+ functions)

Programmable HyperSpeed Engine (PHE)
- Same as BE3 with 1Gb memory
- **PLUS** 32 RISC CPU cores

LINESPEED Integrated Circuits

Networking Signal Management

Retimers, Gearboxes, Mux/Demux for Line Cards and Modules

Gearboxes

100G Gearbox
- with and without RS-FEC

100G Multi-Link Gearbox (MLG)
- 10 x 10GbE Breakout

Retimers

Protocol Independent Retimer
100G (4x25G) Retimer
- with and without RS-FEC

10-Lane Full Duplex 25G Retimer

Mux/Demux

2:1 Serial Multiplexer/Demultiplexor
- Redundant Link Mode option

STELLAR FAMILY of FPGA/ASIC IP

Virtual Accelerator Engines (VAE**) -** ***IP that is Scalable and Portable***
Based on GME (Graph Memory Engine) **Packet Classification Platform**

MoSys Broad Family of Memories & Accelerators

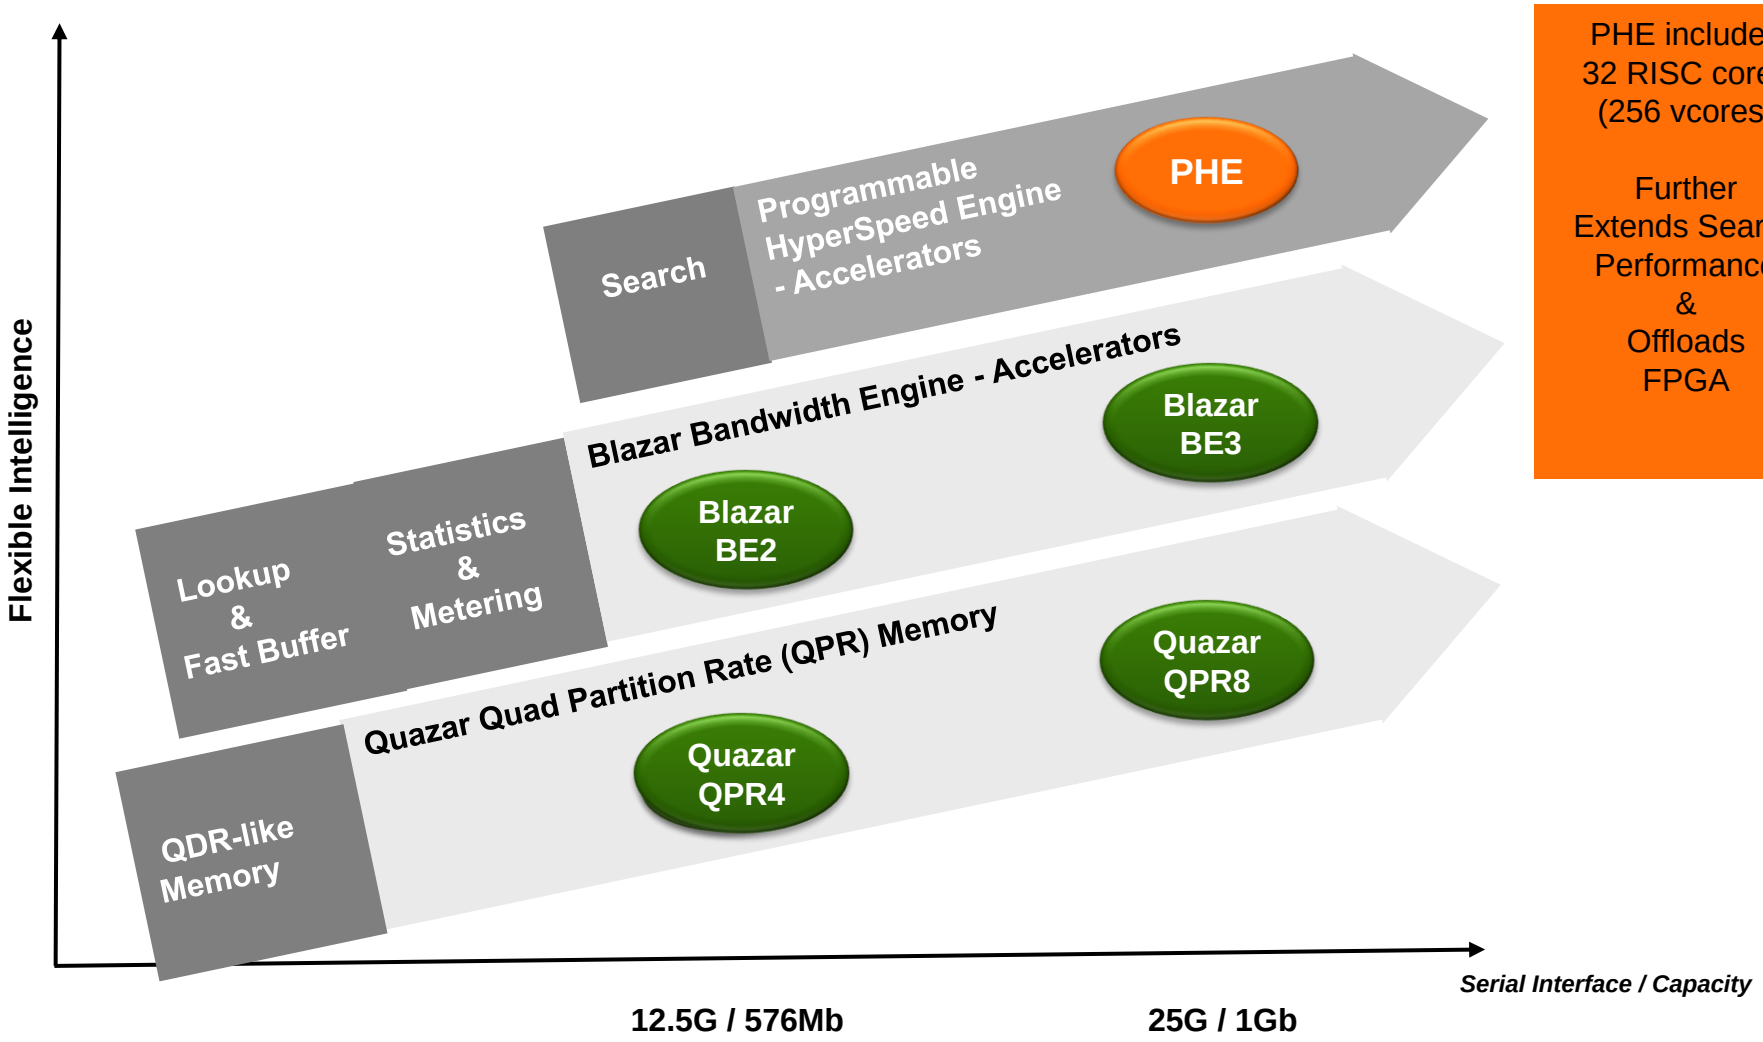

PHE includes
32 RISC cores
(256 vcores)

Further
Extends Search
Performance
&
Offloads
FPGA

Typical Bandwidth Engine Applications

The Bandwidth Engine delivers the Highest Performance of any External Memory

- Most Efficient Interface Protocol
- Up To 25G IO Rates x 16 lanes
- Bandwidth of 640Gbps
- Sixteen concurrent memory ops
- Highest Look-Up Performance
  - ➢ up to 3 billion reads per second

Intelligent Offload – 8 x 100G

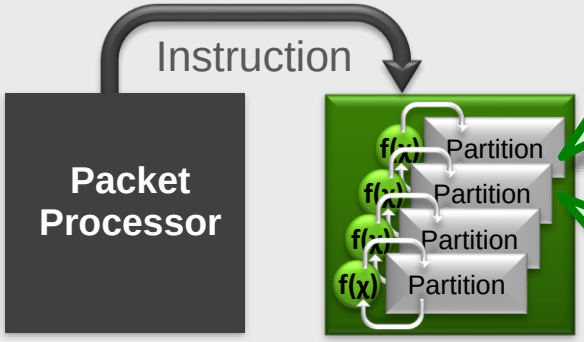

MSR820

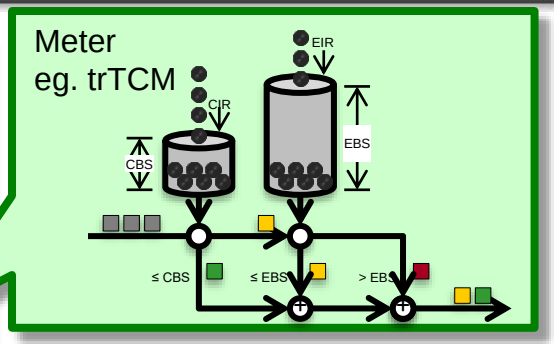

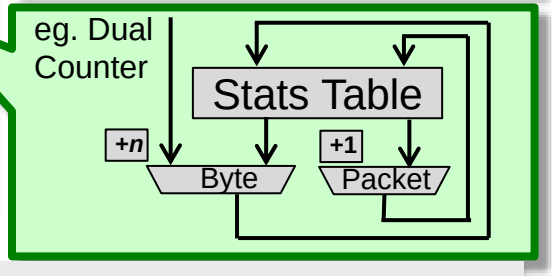

Single Chip - 200 Gbps FDX buffer

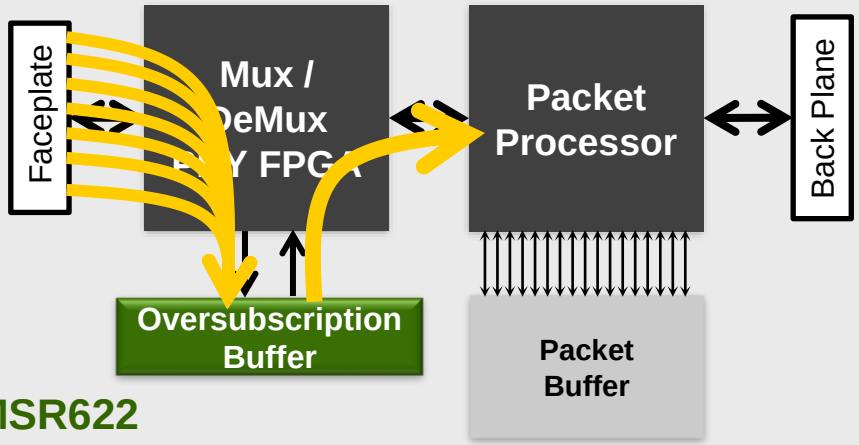

MSR622

High Access Rate Tables– Up to 3B Reads/s

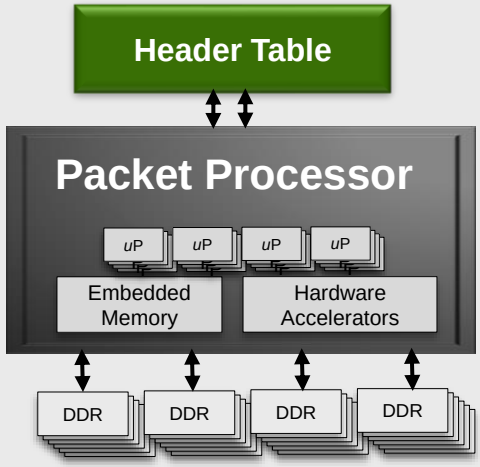

MSR622

Capacity Expansion

- One host port of 16 lanes can connect to 1, 2 or 4 BE devices
- No additional bus loading or pin count
- No throughput degradation

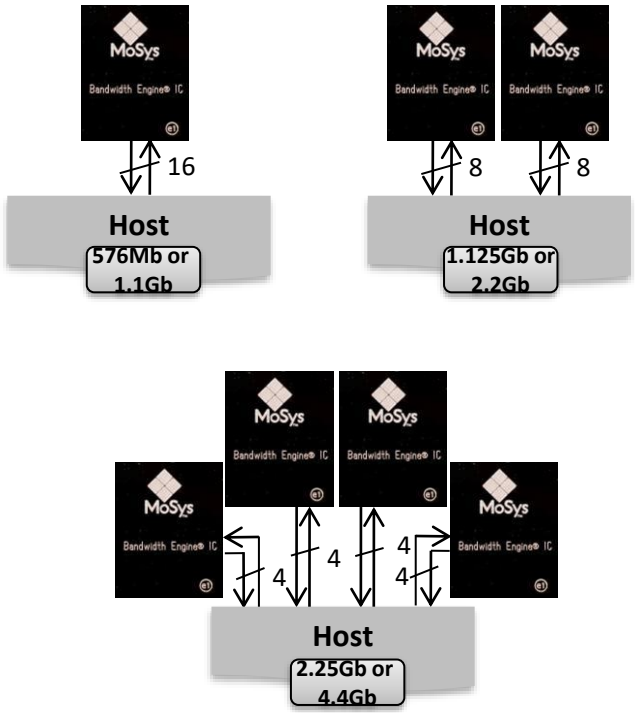

Question 1

- What is your biggest concern when accessing memory?

1)Speed of access
2)Density of Storage
3) Board Space
4) Cost
5) All of the above
6) None of the above

Testing Needs to Achieving Performance

UNDERSTANDING YOUR CHALLENGES

Definitions

- The test and measurement industry helps **manufacturers monitor and improve the quality, safety, health compliance, and productivity of their products**.
- **Electronic test equipment** is used to create signals and capture responses from electronic devices under test (DUTs). In this way, the proper operation of the DUT can be proven or faults in the device can be traced.

Components of Data Gathering

- Electronics measurements
  - Voltage (Voltmeter)
  - Current (Ammeter)
  - Resistance (Ohmmeter)
  - Frequency (Frequency counters)
  - Bandwidth (Function Counters/time)
  - Etc.

Test & Measurement Needs

- Be capable of accurately exercising and measuring the behavior of the DUT
- Be capable of generating the test data patterns to exercise the DUT
- Accept and / or Absorb data at line / device rate
- Absorb bursts of data from a network interface or sampling device(s) like A-D
- Simultaneously received data and transmit / send data to next hop or operation
- Capable of being Calibrated
- Meet all Industry Standards

Test & Measurement Diagram

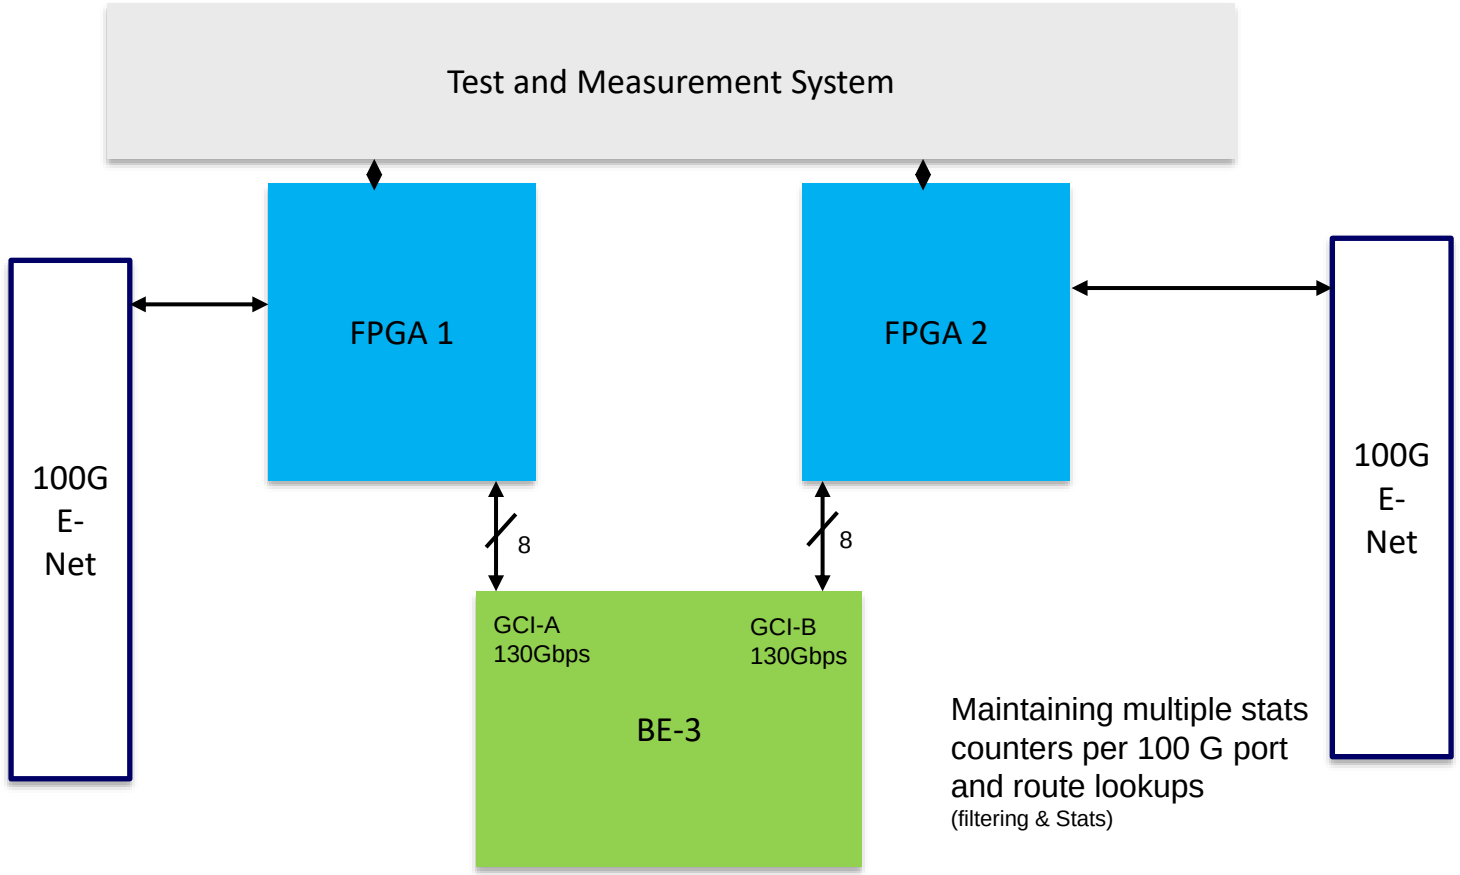

Test & Measurement Diagram (A/D)

32-bit writes are possible in a Blazar

Allows for up to 2.5B writes /sec/port

32M x 32-bit locations

Capable of supporting up to 5B R-M-W / s
- Supports Statistics
- Supports Histograms

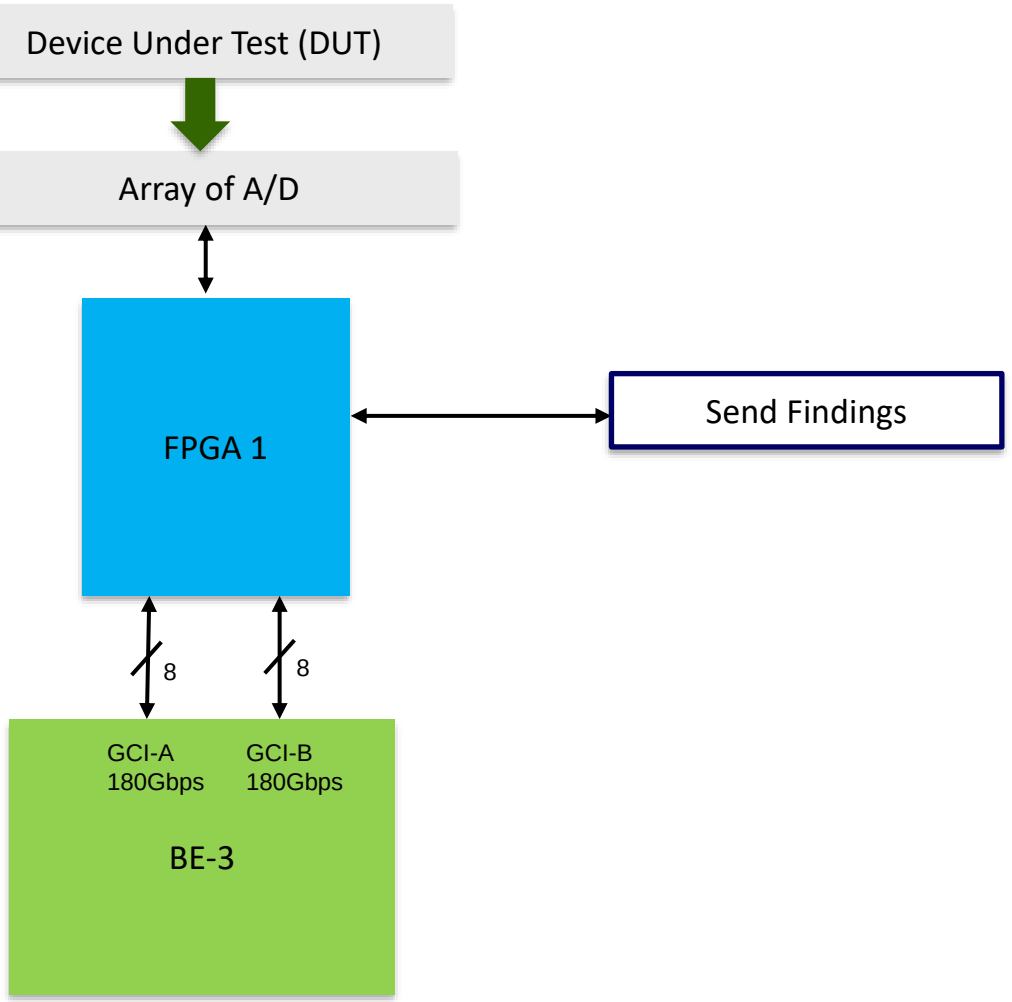

Question 2

- How many parameters are being tested on a common DUT?

1)10’s
2)100’s
3) 1000’s

Accelerator Engine Performance Based on Architecture

QPR4/8
Memory Ops:
Rd/Wr 72b
Wr 36b

BE2:
576Mb, 4 Partitions, 64 Banks
BE3:
1Gb, 4 Partitions, 128 Banks

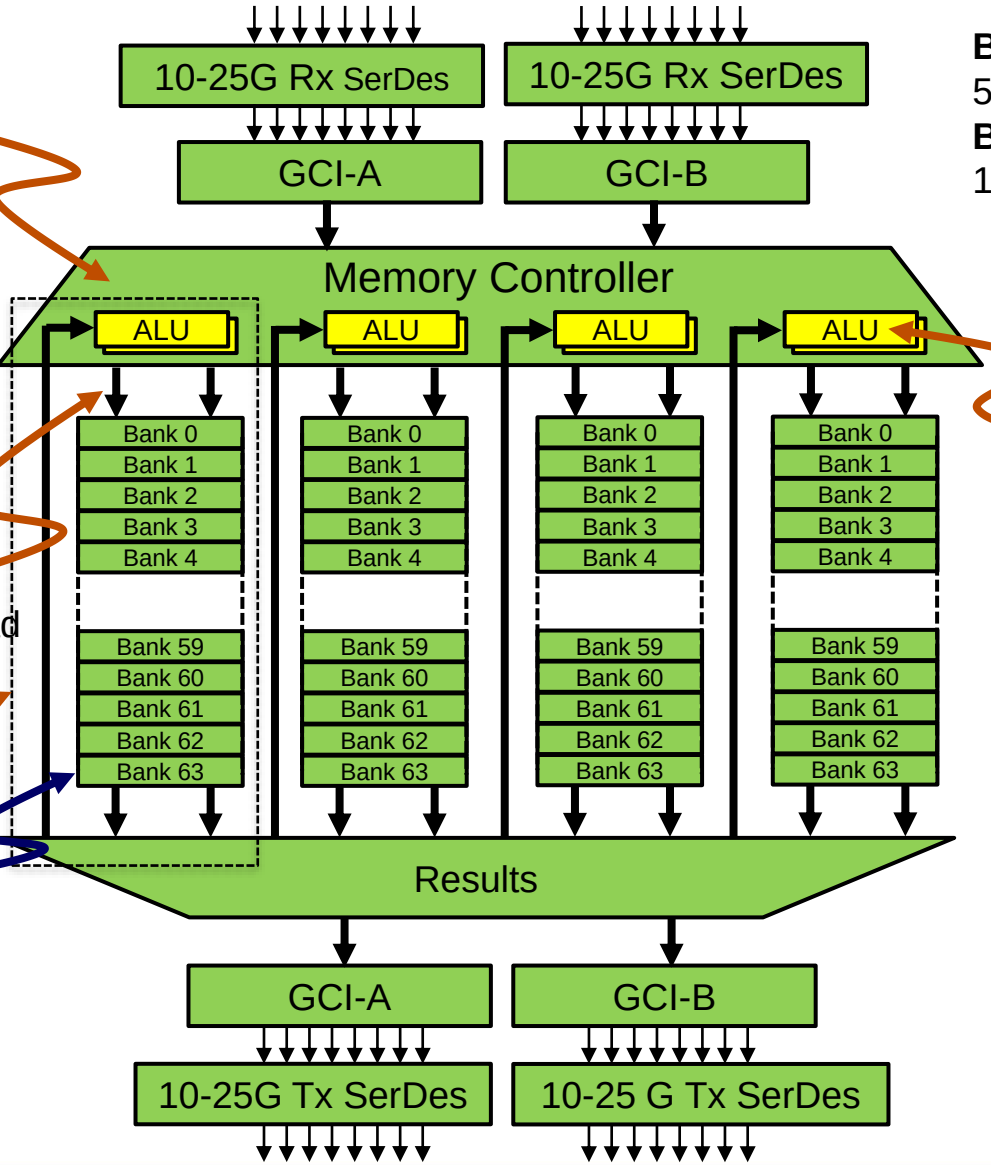

BE2/3 BURST
In Memory Functions
8/16 Reads
8/16 Writes

BE3 - Double the Read & Write capability

Partition

Banks

BE2/3 RMW
In Memory Functions
Incr/Decr
Dual Incr
Metering
Test & Set
Read & Set
More…

4 Partitions Achieve
6B Reads & 6B Writes per second ➔ >1.7 Tbps

4 ports @ 144b
X 4 Partitions

864 Gbps

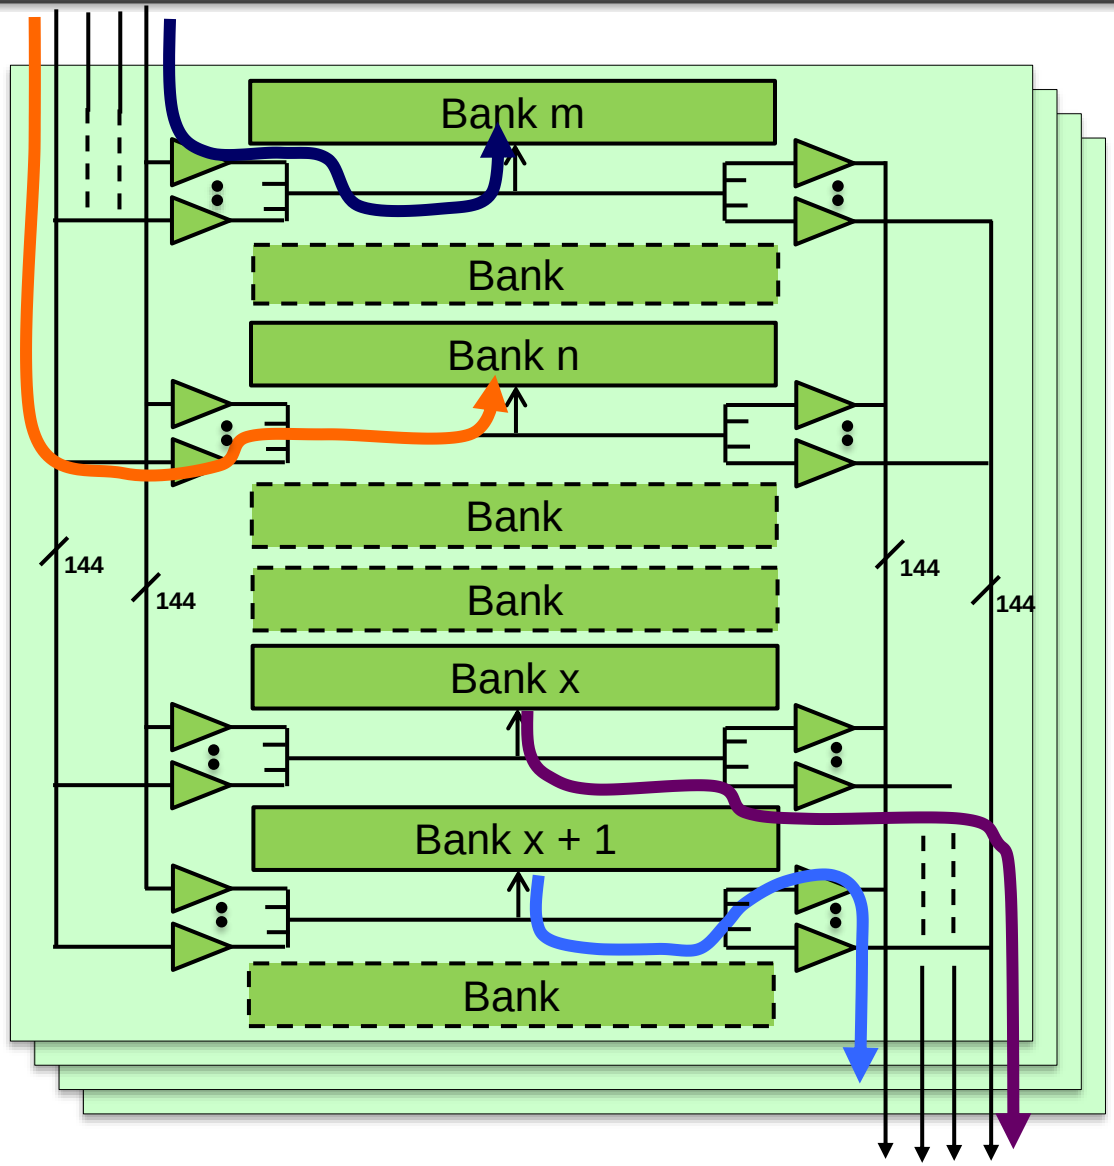

375MHz

4 ports @ 144b
X 4 Partitions

864 Gbps

MSR 622/820/630/830 Single Chip Buffering: Full Duplex Data Throughput

- **Effective throughput of payload; 72b per word**
  - BL# = Burst length; linear burst of 2, 4 or 8 words
- **Full duplex: balanced read and write**

Throughput (Gbps)		Speed Grade			
		622-10	622-12	630-15	630-25
Width	Burst	10.3125G	12.5G	15G	25G
16 lane	BL8	132.0	**160.0**	**192.0**	**320**
	BL4	118.8	144.0	172.8	288
	BL2	99.0	**120.0**	**144.0**	**240**
8 lane	BL8	66.0	**80.0**	**96.0**	**160**
	BL4	59.4	72.0	86.4	144
	BL2	49.5	60.0	72.0	120
4 lane	BL8	33.0	**40.0**	**48.0**	**80**
	BL4	29.7	36.0	43.2	72
	BL2	24.8	30.0	36.0	60

Sigma Quad IVe BL4 is:
93 Gbps Full Duplex
(192 Gbps I/O throughput)
8 x 16Mb single ported banks
~4.5W (device + I/O)

QDR IV XP is:
76.5 Gbps Full Duplex
(153 Gbps I/O throughput)
8 x 16Mb single ported banks
~7W (device + I/O)

RLDRAM 3:
33 Gbps Full Duplex
(76.8 Gbps I/O throughput)
8 x 16Mb single ported banks
~3.5W (device + I/O)

Memory Architecture is Critical
Increase Performance AND Save Space/Cost

Parallel vs Serial

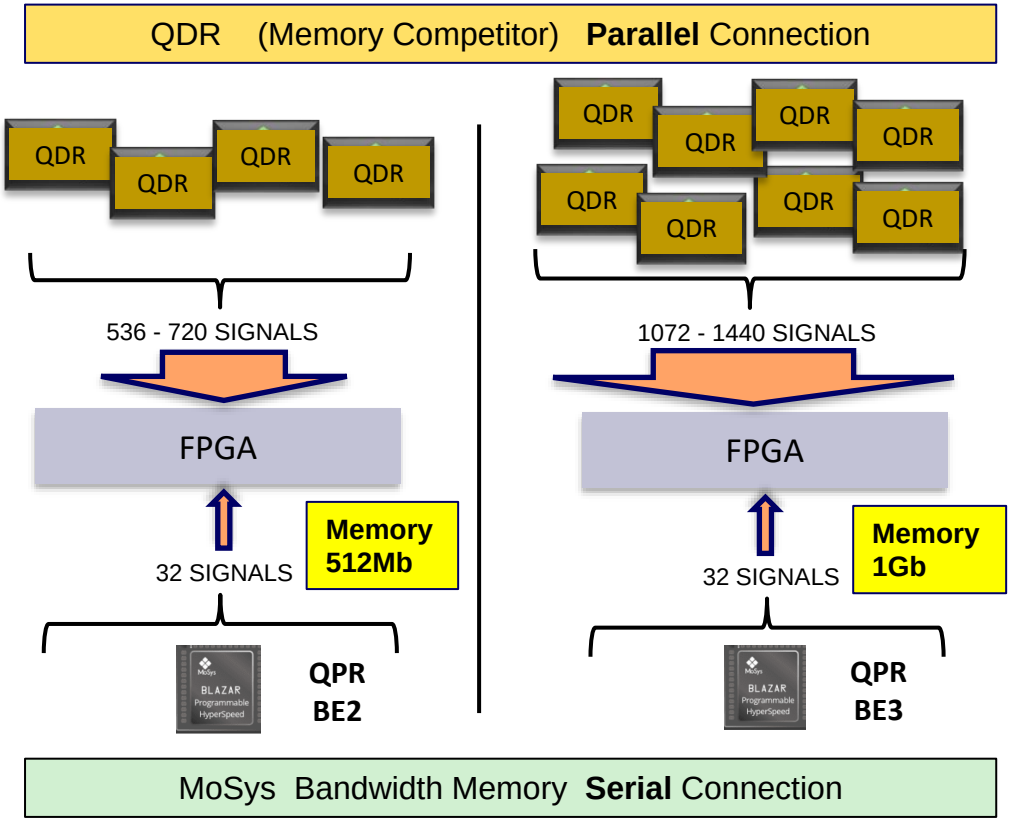

Space & Cost Considerations

Performance Balance of Memory vs Space

- Space and Memory Capacity
  - One QPR/BE2 = 4 QDR … 512Mb
  - One QPR/BE3 = 8 QDR … 1Gb
- Part Cost
  - 1-QPR/BE2 for 512Mb memory ~ 2x cost of one QDR
    - 4x Memory
  - 1-QPR/BE3 for 1Gb memory ~ 2.5x cost of one QDR
    - 8x Memory
- Design time
  - Reduces signal routing time and layout
  - Told it saved 6-9 months
- Signal Integrity
  - Comparable QDR system has 536-1440 clean signals generally requiring external components
  - MoSys system typical has 32 signals with on board Auto-Adaptation signal tuning
    - No external components
- Power
  - ~ Half
- Bandwidth
  - Random data access is equivalent (tRC 2.7-3.2 ns)
  - For certain applications, much faster

Performance impact of In-Memory Functions BURST-RMW

LET THE MEMORY DO THE WORK!

BE3 - BURST

MSR630 – Bandwidth Engine IC

- Fast, Efficient Serial Interface
  - 16 lanes, 10-15G, 20-28G SerDes
  - Independent GCI Ports
- Highest Performance Single Chip Transaction Rate Memory
  - 1Gbit memory capacity
  - 5+ billion transactions/sec
  - 32 ported memory array architecture
  - 3ns memory cycle time
- Burst Mode Capable
  - Burst of 2, 4, or 8
- Testing modes
  - PRBS generation and checking
  - Loopbacks
- Small form factor
  - 27x27mm package

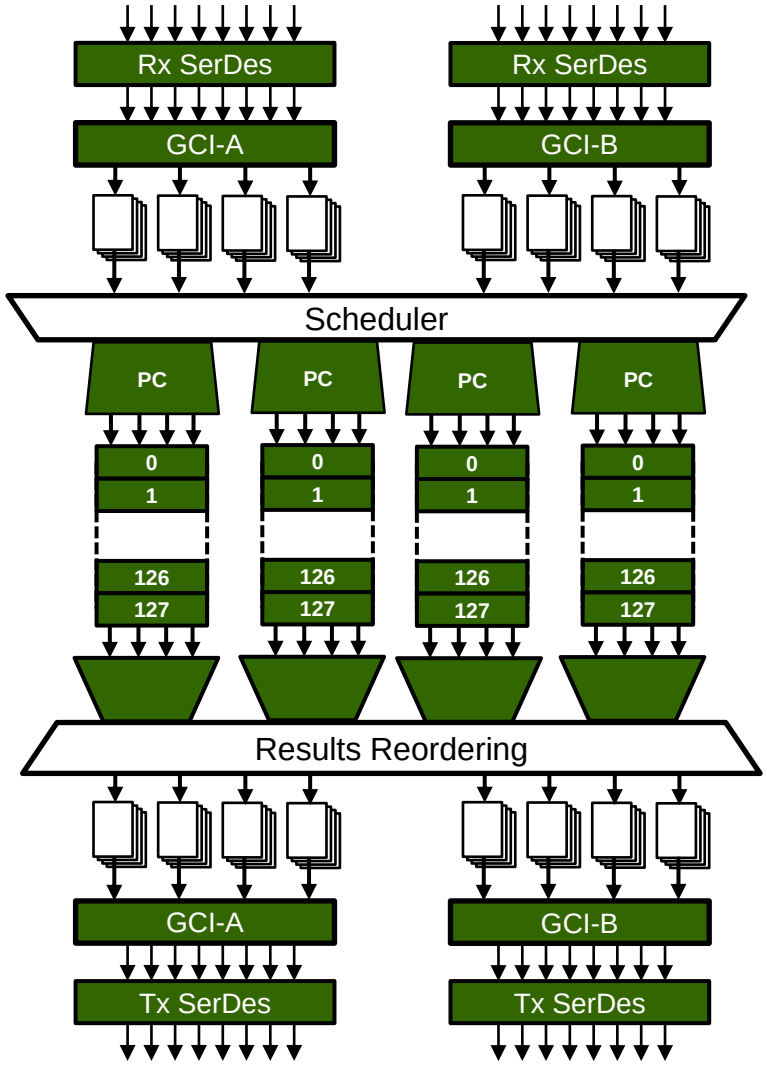

BE3 – RMW

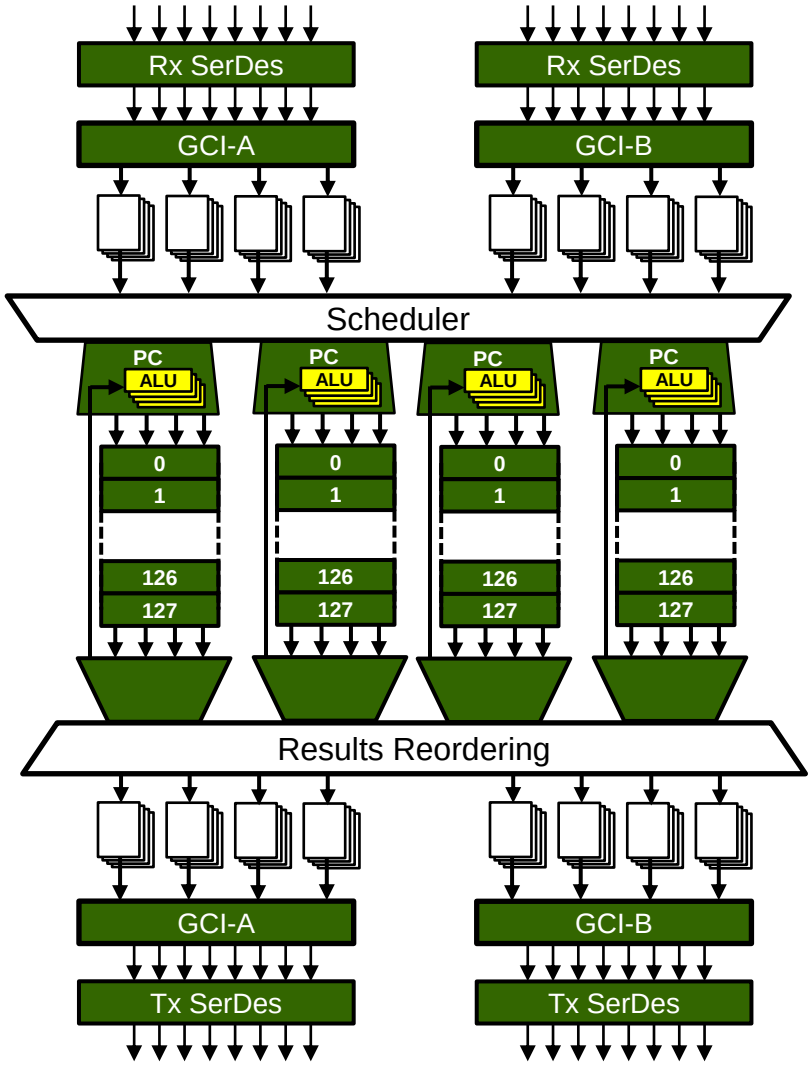

BE with Intelligent Offload

- Fast, Efficient Serial Interface
  - 16 Lanes, 10-15 or 20-28Gb/s
  - Independent GCI Ports
- Intelligent Macro Offload
  - Multiple ALUs per Partition
  - Built in Memory coherent statistics, Metering and atomic operations
- Macro optimize transaction efficiency
  - Low pin count but high performance
- Highest Performance Single Chip Transaction Rate Memory
  - 1Gbit memory capacity
  - 5+ billion transactions/sec
  - 32 ported memory array architecture
  - 3ns memory cycle time
- Burst Mode Capable
- Small form factor
  - 27x27mm package

New Counter Capabilities

- **Single Instruction 'Dual' counter for per flow packet & byte**
  - Single command updates two counters
  - 16b immediate (byte #) to byte counter, Inc/Dec the packet counter
  - Operates on two 64b "Lifetime" counters

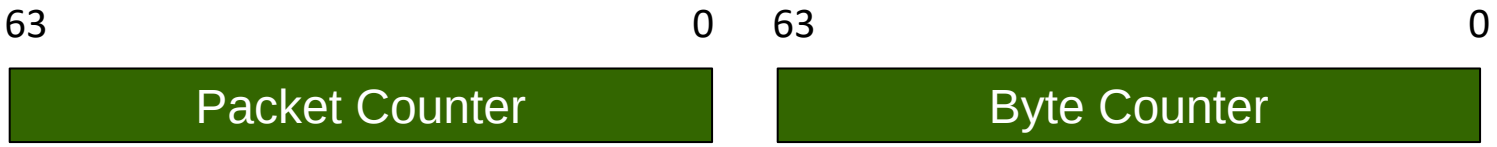

- **Single Instruction Dual 'Split' counter**
  - User selectable division of a 64b value
  - Single command performs dual counter function, same as above

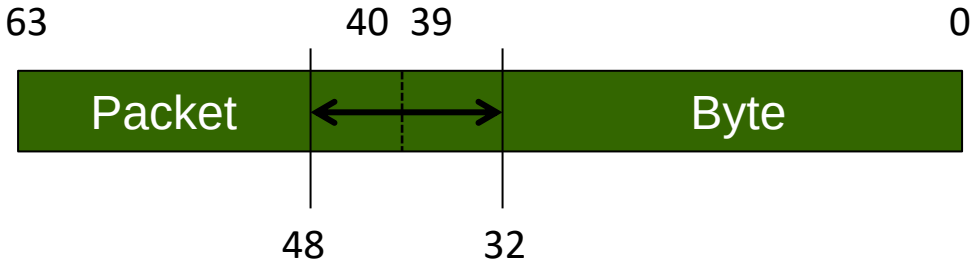

Pipelining With BE3 As A Dual Port

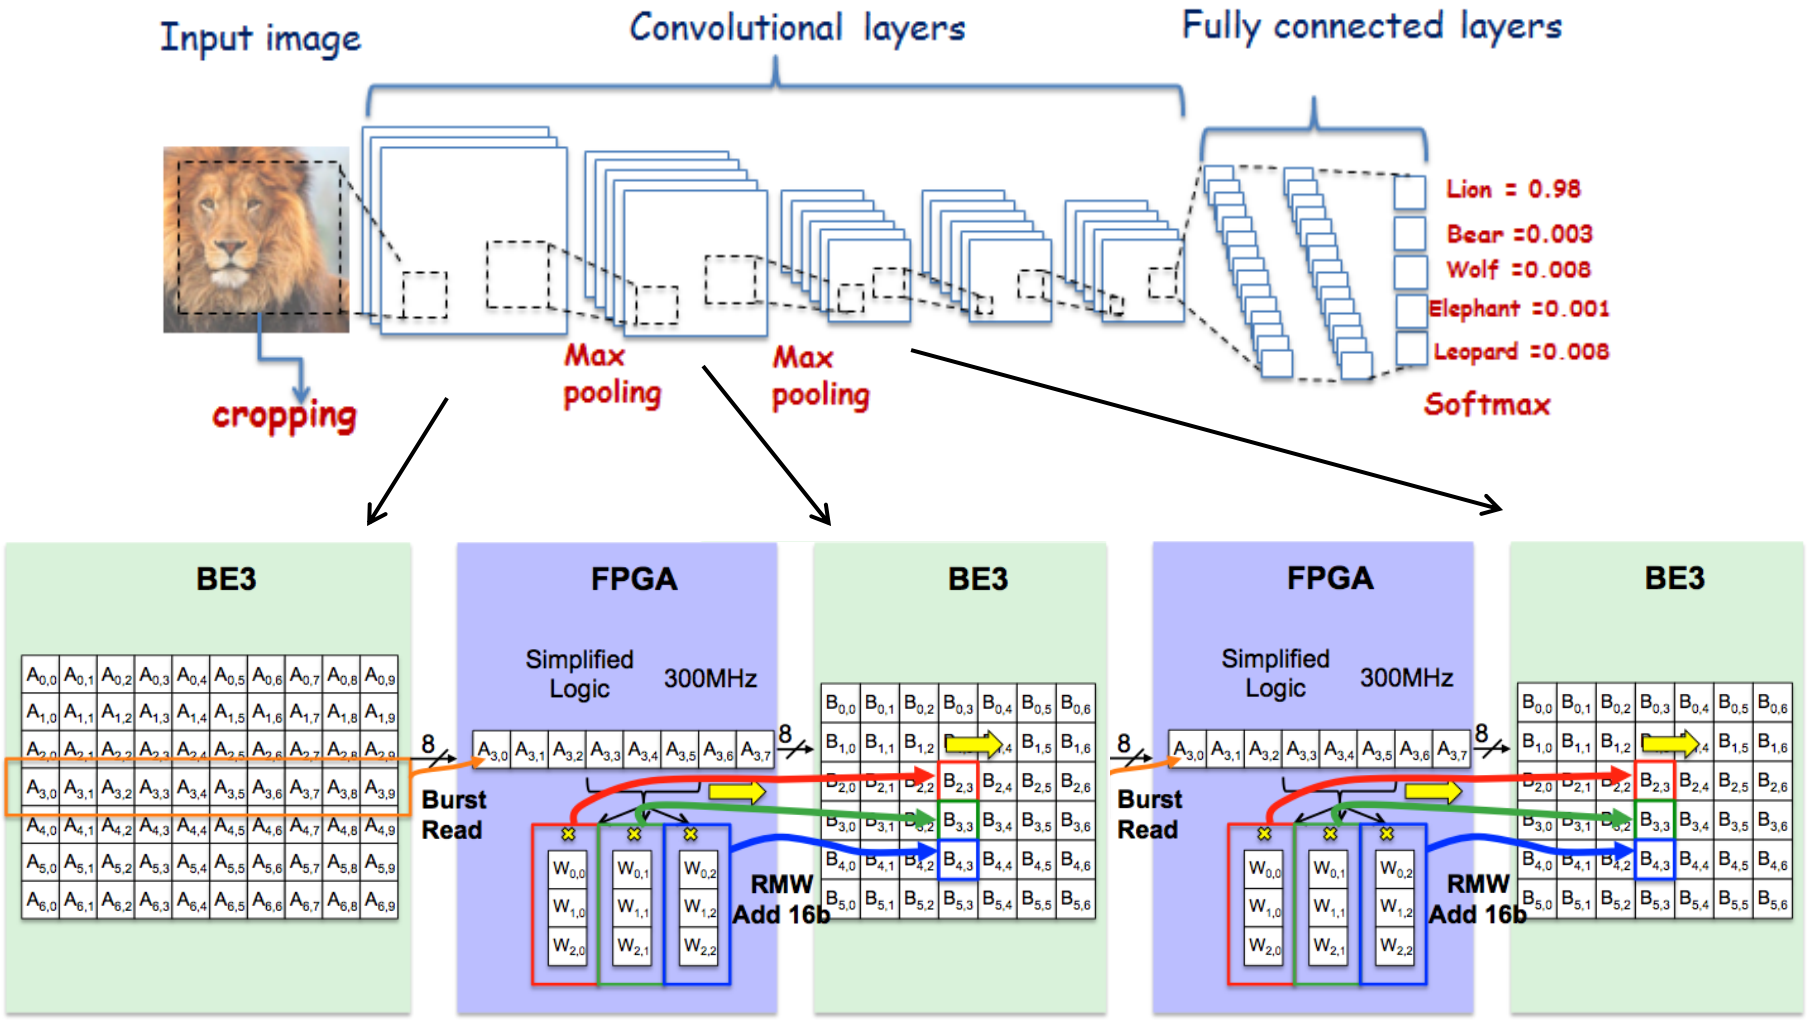

Question 3

- Does your system utilize high speed SRAM (QDR) for high-speed memory access?

1)Yes
2)No
3) Yes but a larger device would be very beneficial to the design.

Programmable Hyper-Speed Engine

Programmable HyperSpeed Engine (PHE) Architecture

- **Physical**
  - 16 x 10 to 25Gbps PHY
  - SerDes Standard GigaChip Interface (GCI) Protocol
  - 8 Scheduling Domains
  - Integrated 1Gb Fast Memory
- **Threaded Processor Engines**
  - 32 cores (8 clusters of 4 PE)
  - Up to 1.5GHz
  - 8 threads per PE
  - Search-optimized ISA
- **Internal Performance**
  - 5B Rd/s + 5B Wr/s
  - 3ns tRC
  - Access up to 144 bits/cycle
- **676 FCBGA 27x27mm, 1mm**

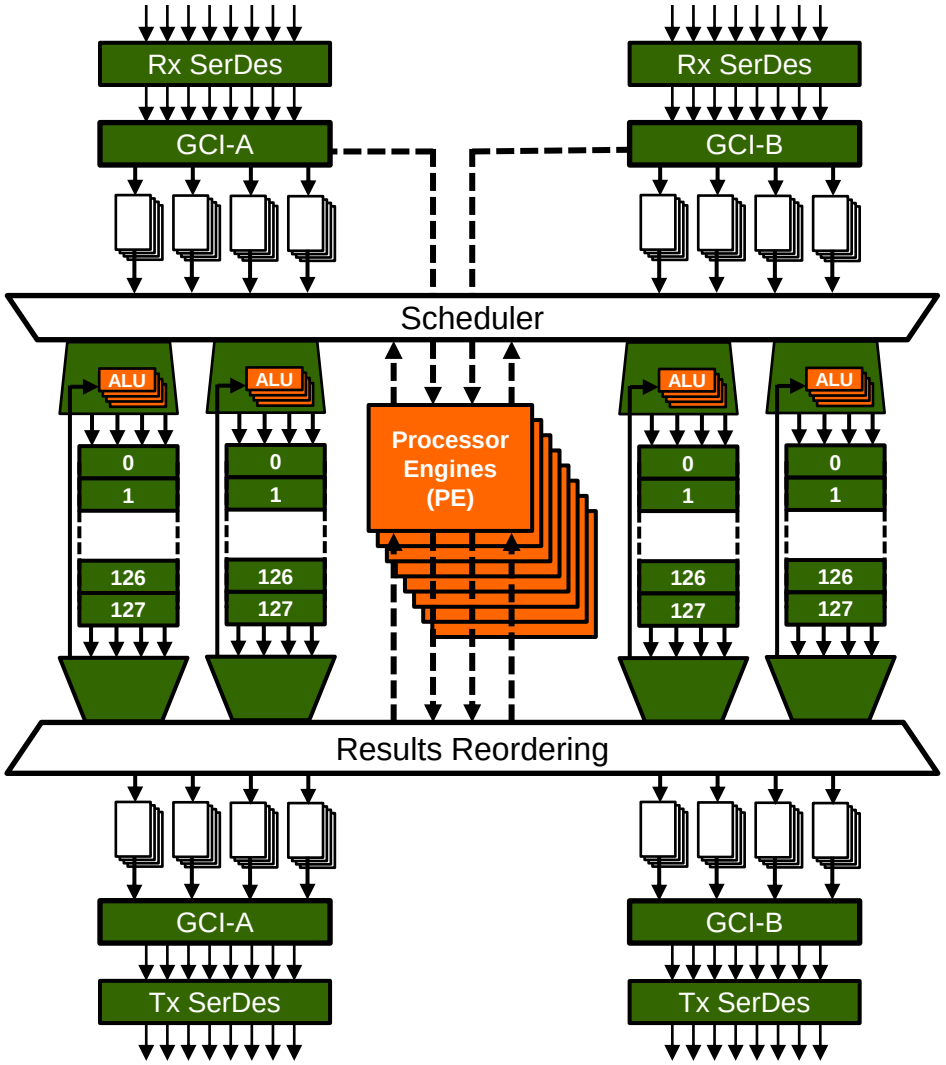

Instruction Set Overview

- **ALU/Logical on 72b**
  - add, sub, adc, sbb, s1add, s2add, s3add, s3sub, and, or, xor, andn, sar, slr, sll, minu, maxu, mult
- **Bit field of variable len @ variable pos**
  - Extract, deposit, chomp
  - Can be across register boundaries
  - Optional auto incr of pos
- **Special Functions**
  - Find first zero, find first one
  - Population count
  - Swap bits in bytes and bytes in words
  - 144b HASH to 72b (non-crypto)
  - Compute CRC32
  - Mult-way compare with 4, 6, 9 & 12 inputs
- **Test & Branch**
  - tsteq, tstgt, tstnle, tstlt, tstnge, tstgtu, tstnleu, tstltu, tstngeu, tstbs, tstne, tstle, tstngt, tstge, tstnlt, tstleu, tstngtu, tstgeu, tstnltu, or tstbc
  - Jmp, jeq, jgt, jnle, jlt, jnge, jgtu, jnleu, jltu, jngeu, jbs, jne, jle, jngt, jge, jnlt, jleu, jngtu, jgeu, jnltu, or jbc
  - Multiway branch 2, 3 & 4
- **Loads & Stores**
  - Local Dmem:
    - 8, 16, 32, 64 & 72b
    - Reg + offset, w/auto incr reg
  - Partition:
    - Burst reads, load balanced reads and broadcast
    - 64, 72, 128, 135, 144b
    - Reg + reg or reg + offset, w/auto incr reg
- **Atomic Operations**
  - Local:
    - 8, 16, 32 & 64b
    - adda, suba, anda, xora, andna, xchga, cmpxchga
  - Partition:
    - 16, 32, & 64b
    - Add(s), sub(s), xor, rd/set, tst/set, cmp/set, avg, tm, age
- **Program Control**
  - Hlt, Brk & nop
  - Add/mov & halt (tread)
  - Yield
- **Special registers**
  - GPR indirect specification
  - Auto increment
  - Command, memory, result, result len
  - Time stamp, random, zero, all ones, thread id, wake up, sink

Accelerator Engine Devices

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
			mm	Tx/Rx	10.3	12.5	15.6	25	Gb	ns	Gb				
QPR4	MSP220	QPR4 (Quad Partition Rate) 0.5 Gb	FCBGA 19X19	16	✓	✓			120	3.2	O.5	2.5	✓		
QPR8	MSP230	QPR8 (Quad Partition Rate) 1Gb	FCBGA 27X27	16			✓	✓	240	3.2	1	4	✓		
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 QPR (Quad Partition Rate) Memories	FPGA RTL Code		✓	✓	✓	✓			576Mb & 1Gb	6.5	✓	✓	✓

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

Thank You