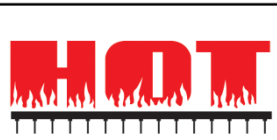

Hot Chips 25
Stanford Memorial Auditorium
August 25-27, 2013

Second Generation Bandwidth Engine® IC Breaks 4.5 Billion Accesses/sec

Michael J. Miller
VP, Technology Innovation & System Applications

The Vision: Fast, Intelligent Access Architecture

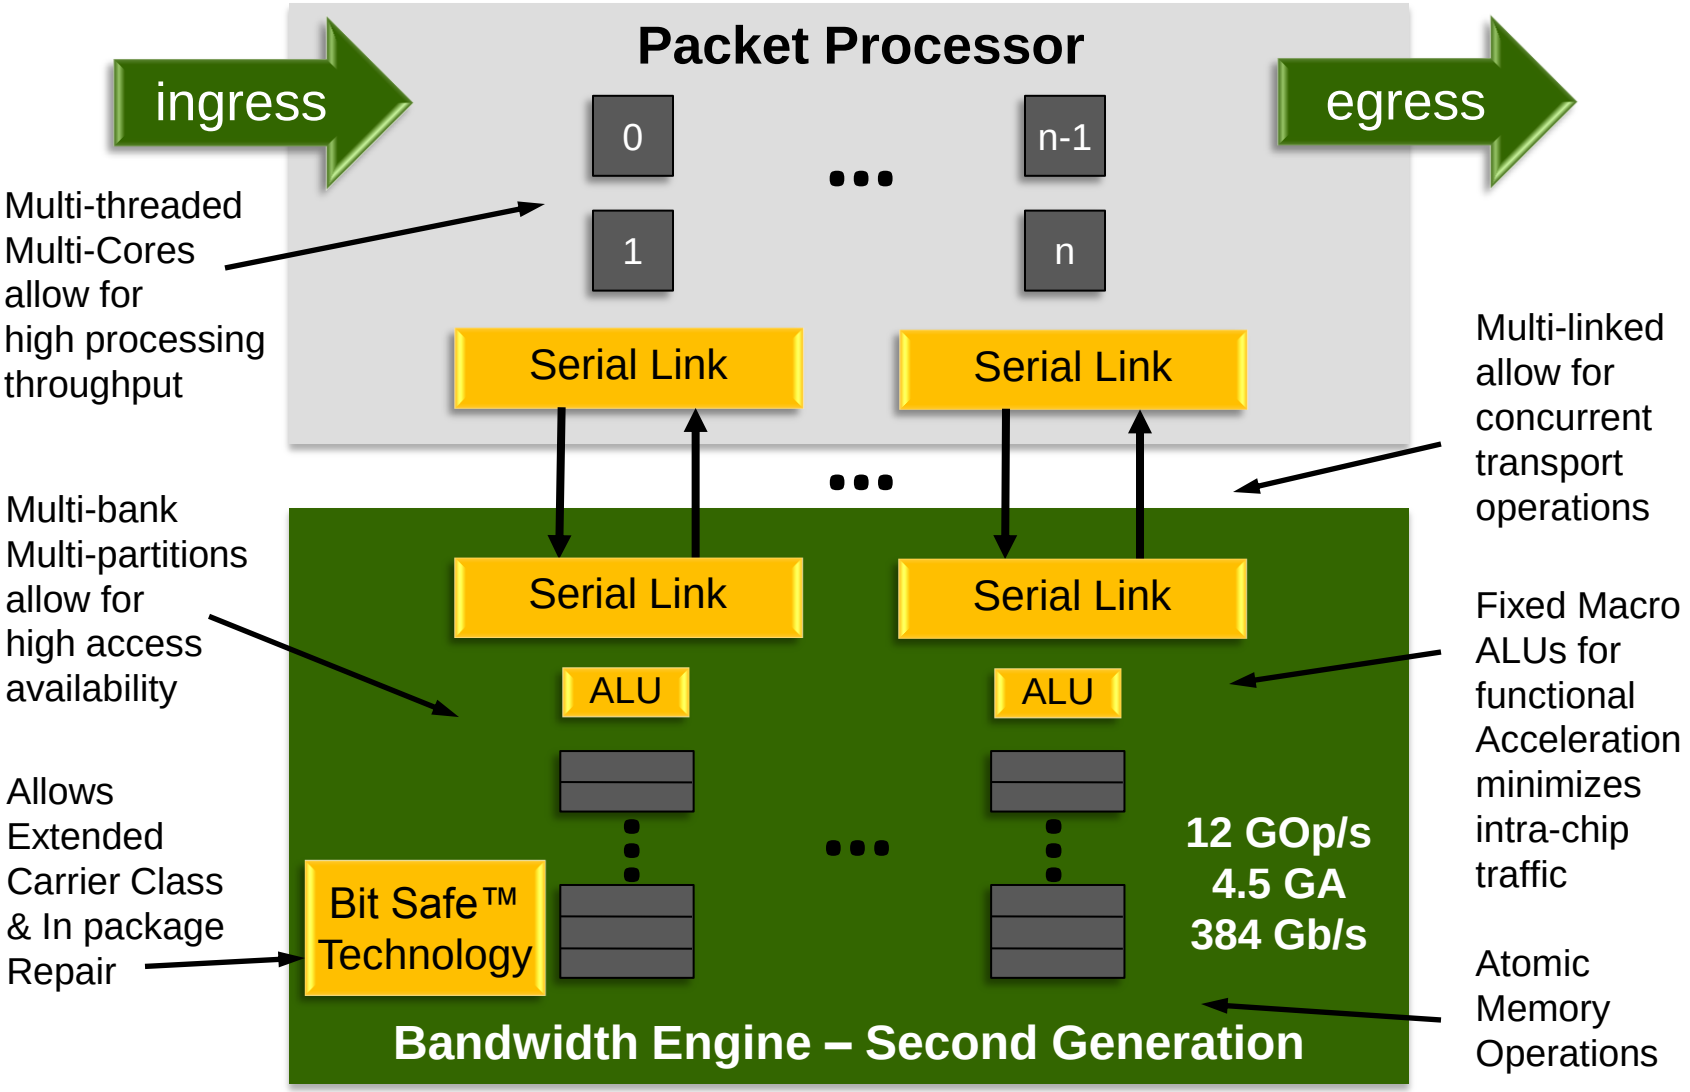

Bandwidth Engine 2 Architecture & Family Sampling Now

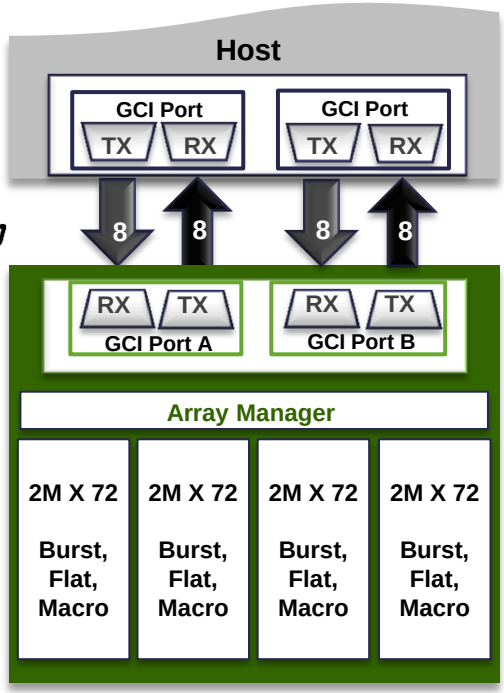

- **Parallel Array Architecture … Performance up to:**
  - 16 outstanding transactions
  - 4.5G Accesses per second
  - 192 Gbps full duplex throughput
  - ~12ns deterministic read latency
  - 2.7ns Random cycle time (tRC)
- **GigaChip Interface … 90% Efficient Transport Protocol**
  - Up to sixteen low latency SerDes lanes (8G to 15G)
- **High Reliability …70X better SER than 6T-SRAM**
  - Full ECC support; 72bit array and macro datapath operations
  - CRC protected and self recovering GigaChip Interface
  - SEU resistant 1T-SRAM Memory core: < 10 FIT/Mb
  - Bit Safe™ Self Test and Self Repair Option

	Bandwidth Engine – 2nd Generation			
Applications	**BE1 - MSR576**	**MSR620**	**MSR820**	**MSR720**
Lookup, LPM, Hash	Highest single component Acce s Rate			
Statistics	Onboard ALU		Onboard ALU+	
Buffer – up to 80% efficient		Per cycle Burst, Write Broadcast		
Metering, Dual Ops			Fixed Macro	
Semaphore, Link List			Atomic Ops	
State, Queuing, Link List				Dual Port w/ Data Coherency, 36 bit word access

Functional Design

Bandwidth Engine MSR720 Architecture

Memory Ops:
Rd/Wr 72b
Wr 36b

Bank Conflict Resolution:
32K x 72b SRAM
Delayed Write Cache

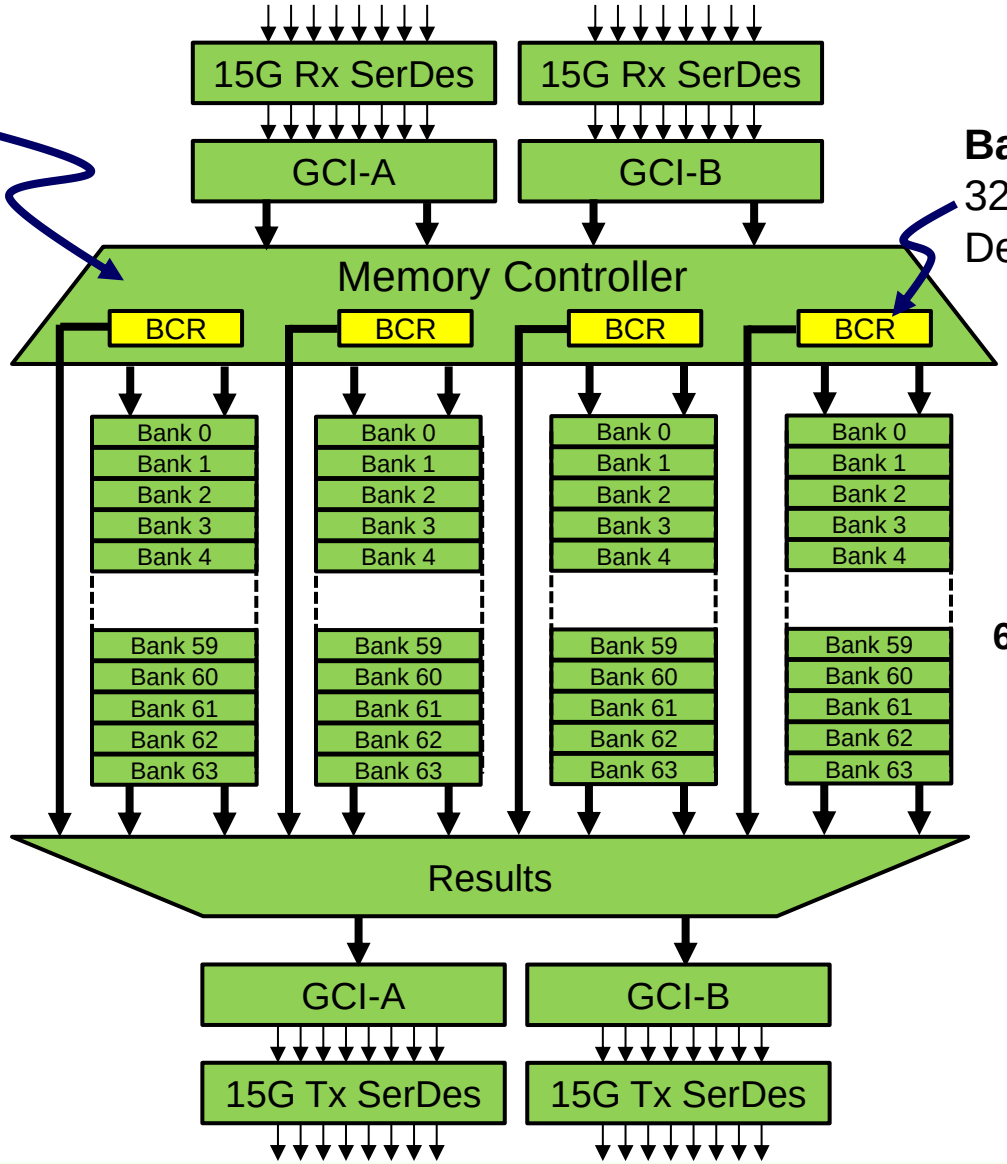

375 MHz Clock x
8 Reads +
8 Writes
➔ 6GA Internally

4 x Partitions:
64 Single Port Banks
32K x 72b each

MSR720: Bandwidth Engine 2 – Access

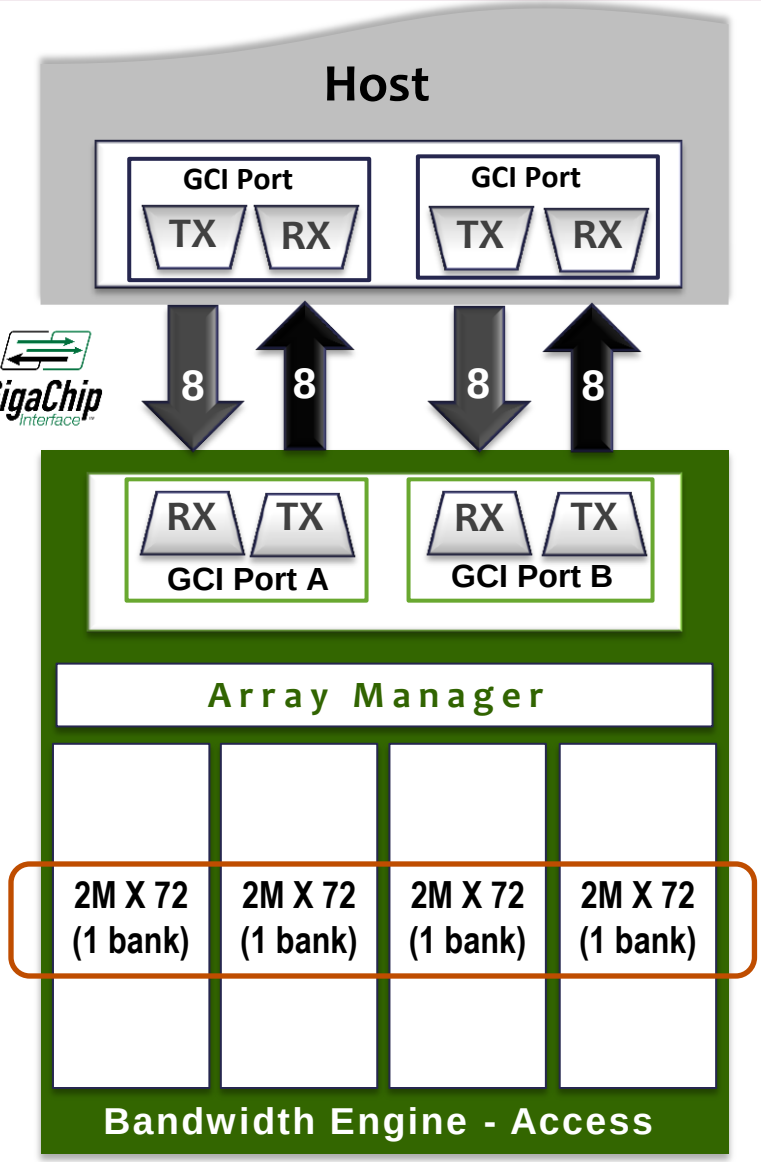

- **Simultaneous Read & Write @ same address**
  - **… treat each partition as a single bank**
  - **… >>2x the access performance of QDR SRAM**
    - 4.5GA : 3 billion reads/sec, 1.5B 36b write/sec
    - 72b & **36b** words each access
    - 12 ns read latency pin to pin
    - 2.7ns tRC cycle time
    - 8.5W @ 12.5G system power
  - **… 90% efficient transport protocol**
    - Up to Sixteen 15G serial lanes (2 links of 8 lanes)
- **High Density**
  - **… 576Mbit 1T-SRAM ® memory core**
    - 19mm x 19mm package – 1mm pitch
- **High Reliability**
  - **… 70X better SER than 6T-based SRAM**
    - Memory core: < 10 FIT/Mb native
    - Interface: < 1 FIT
    - Bit Safe™ Self Test and Self Repair Option

MSR720: Basic 'Dual Port' Operation

Flat Partition Mode:

- Double the effective bandwidth in worst case:
  - Bank Conflict Resolution (BCR) function allows for simultaneous Read and Write of the same bank
  - Implements dual port behavior with single port banks
- Guarantees data coherency
- 3 billion accesses / sec @ 15G

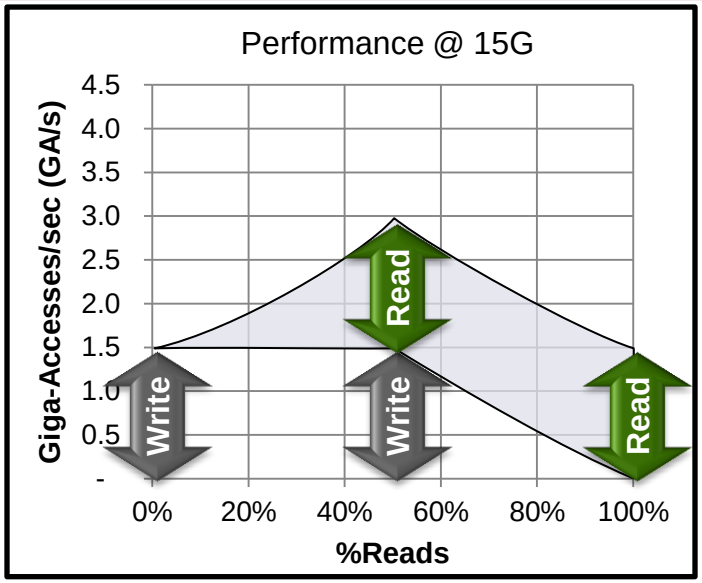

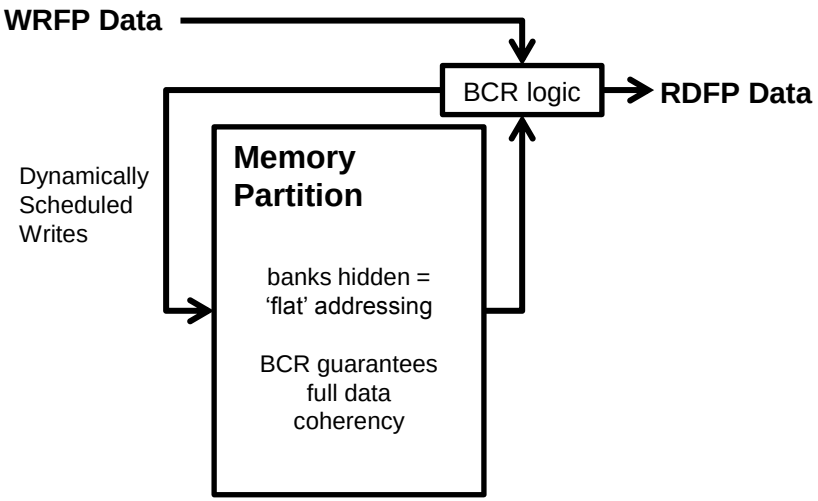

Scheduling Balanced R:W

Balanced R:W Controller

Frame	Partition	Input (RX) CMDARX		Input (RX) CMDBRX		Frame	Output (TX) QATX	Output (TX) QATX
1	0/1	RDFP	WRFP	RDFP	WRFP	16	RDFP	RDFP
2		WD		WD		17		
3	2/3	RDFP	WRFP	RDFP	WRFP	18	RDFP	RDFP
4		WD		WD		19		

15G

1.50 GA Write
1.50 GA Read
3.00 GA Total

"BCR" Implemented w/Delayed Write Cache Conceptual Block Diagram

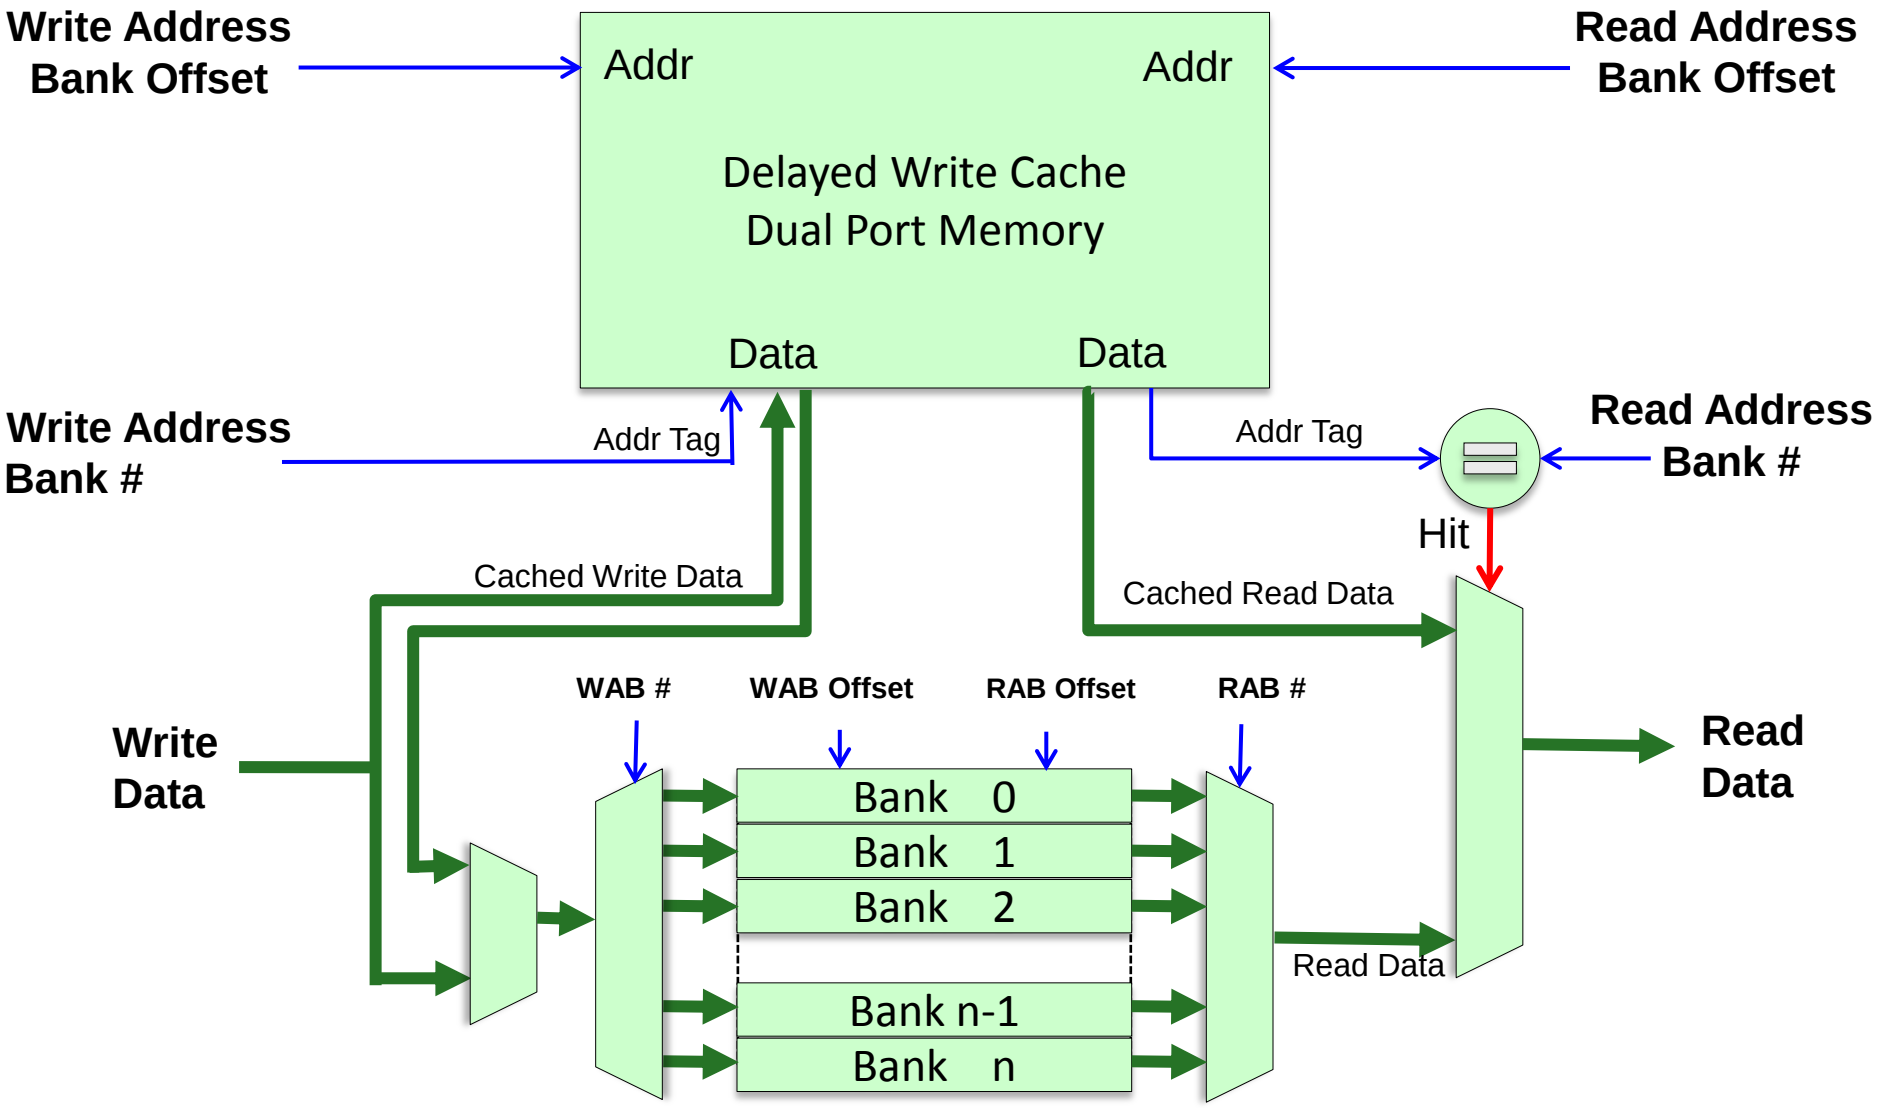

Up to 4 Accesses Per Partition In One Cycle
3 Accesses Sustained Throughput

- **The MSR720 supports Bank Aware commands.**
  - This can be used to improve the data read performance on the interface, however bypasses the BCR logic

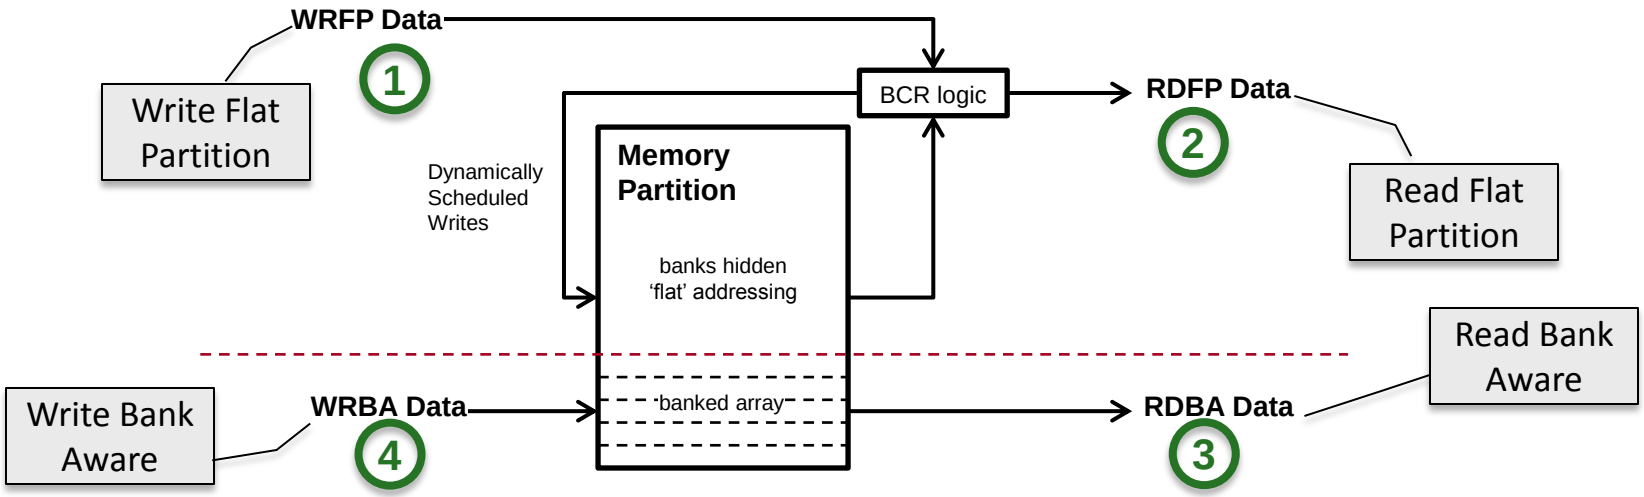

  - Useful for unified memory applications combining dual port SRAM performance of "buffers" or "state" tables and read-only "lookup" tables
  - The two address ranges cannot overlap.

- **MSR720 supports 36b write operations: WRFP and WRBA**
  - Useful for small word size tables (pointers)
  - Increases write performance

Mapping Ports and Banks

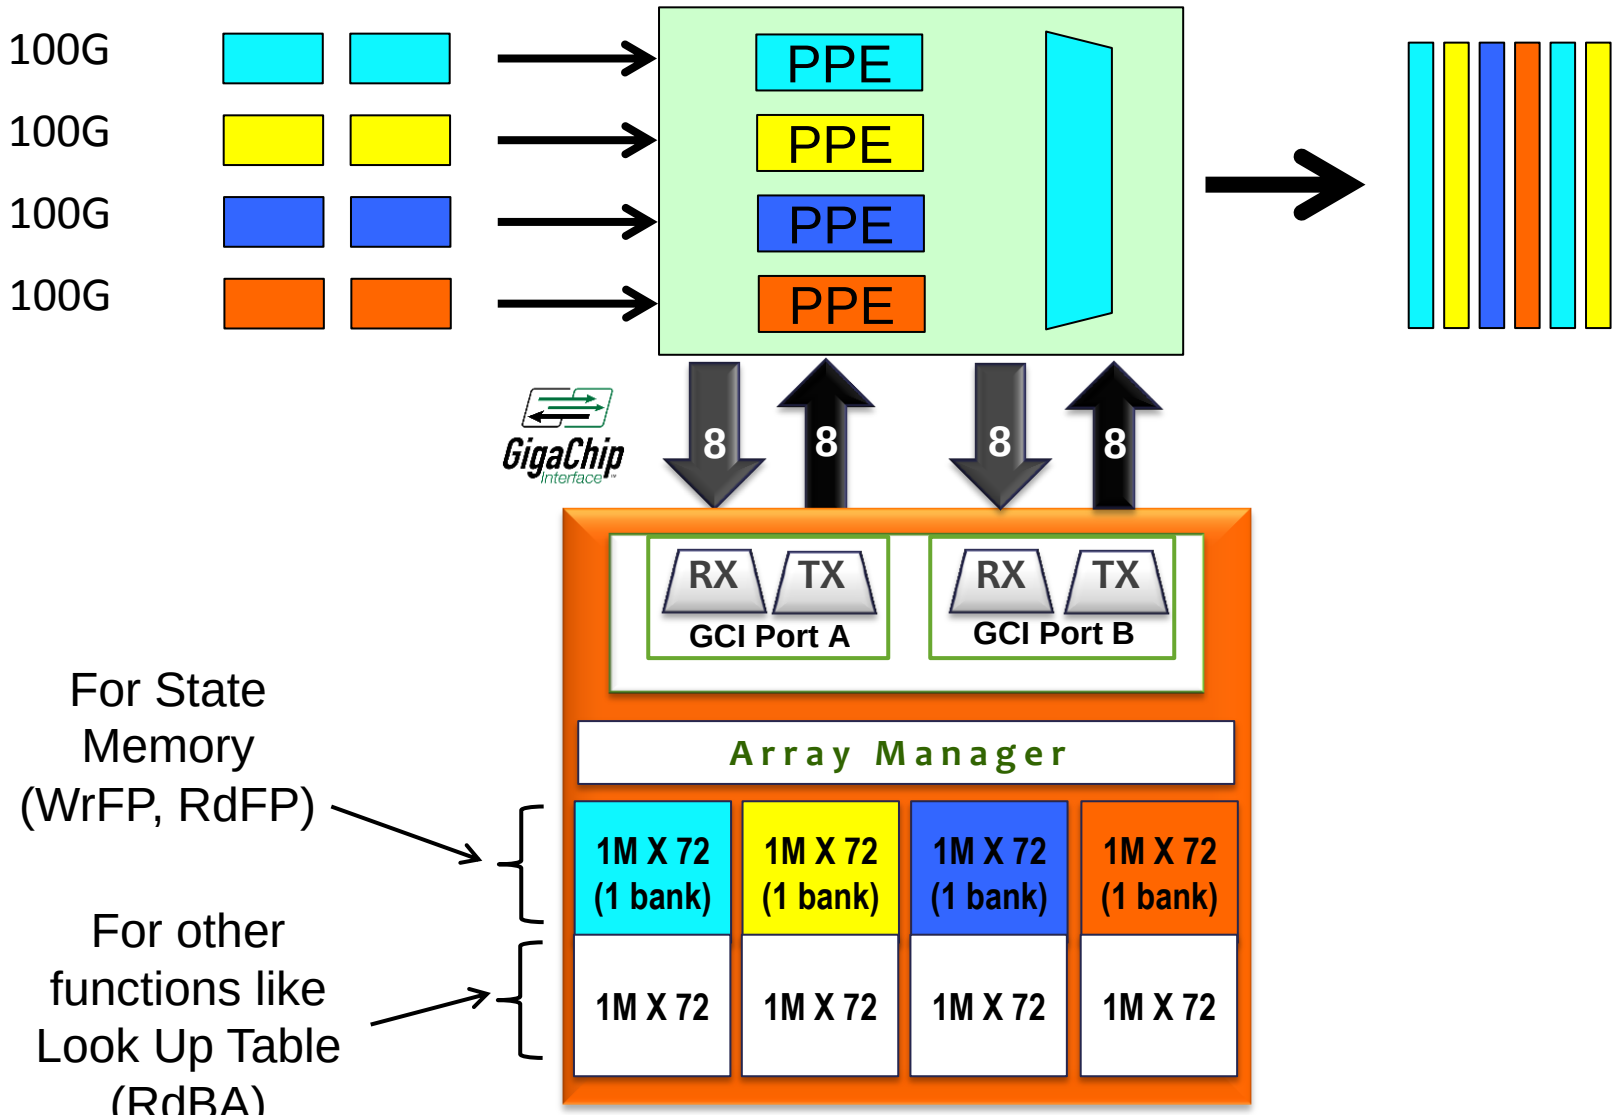

MSR720 Access Scheduling

- **Flat Partition guarantees no bank conflict between read and write to any address, even in the same cycle**

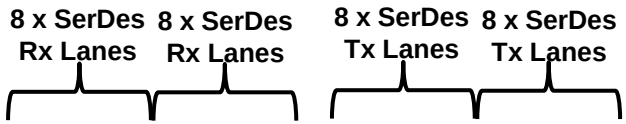

2.7ns

Balanced R:W Controller (Basic 72b R:W Operation)

Frame	Partition	Input (RX) CMDARX		CMDBRX		Frame	Output (TX) QATX	QATX
1	0/1	RDFP	WRFP	RDFP	WRFP	16	RDFP Data	RDFP Data
2		WD		WD		17		
3	2/3	RDFP	WRFP	RDFP	WRFP	18	RDFP Data	RDFP Data
4		WD		WD		19		

15G

1.50 GA Write **FP**
1.50 GA Read **FP**
3.00 GA Total

WD = 72b data

Native BE Controller (72b WRITES)

Frame	Partition	Input (RX) CMDARX		CMDBRX		Frame	Output (TX) QATX	QATX
1	0/1	RDFP	**RDBA**	RDFP	**RDBA**	16	RDFP	RDFP
2		WRFP	WDL	WRFP	WDL	17	**RDBA**	**RDBA**
3	2/3	RDFP	**RDBA**	RDFP	**RDBA**	18	RDFP	RDFP
4		WRFP	WDL	WRFP	WDL	19	**RDBA**	**RDBA**
1	0/1	RDFP	**RDBA**	RDFP	**RDBA**	16	RDFP	RDFP
2		WRFP	WDU	WRFP	WDU	17	**RDBA**	**RDBA**
3	2/3	RDFP	**RDBA**	RDFP	**RDBA**	18	RDFP	RDFP
4		WRFP	WDU	WRFP	WDU	19	**RDBA**	**RDBA**

15G

1.50 GA Read **BA**
0.75 GA Write **FP**
1.50 GA Read **FP**
3.75 GA Total

WDL = Lower 36b data
WDU = Upper 36b data

Native BE Controller (36b WRITES)

Frame	Partition	Input (RX) CMDARX		CMDBRX		Frame	Output (TX) QATX	QATX
1	0/1	RDFP	**RDBA**	RDFP	**RDBA**	16	RDFP	RDFP
2		WRFP	WD	WRFP	WD	17	**RDBA**	**RDBA**
3	2/3	RDFP	**RDBA**	RDFP	**RDBA**	18	RDFP	RDFP
4		WRFP	WD	WRFP	WD	19	**RDBA**	**RDBA**

15G

1.50 GA Read **BA**
1.50 GA Write (36b) **FP**
1.50 GA Read **FP**
4.50 GA Total

Partition Access Restrictions

Frame	Partition	RX GCI-A	GCI-B
1	0/1	P0	P1
2			
3	2/3	P2	P3
4			

MSR 720 : Breaking 4.5GA

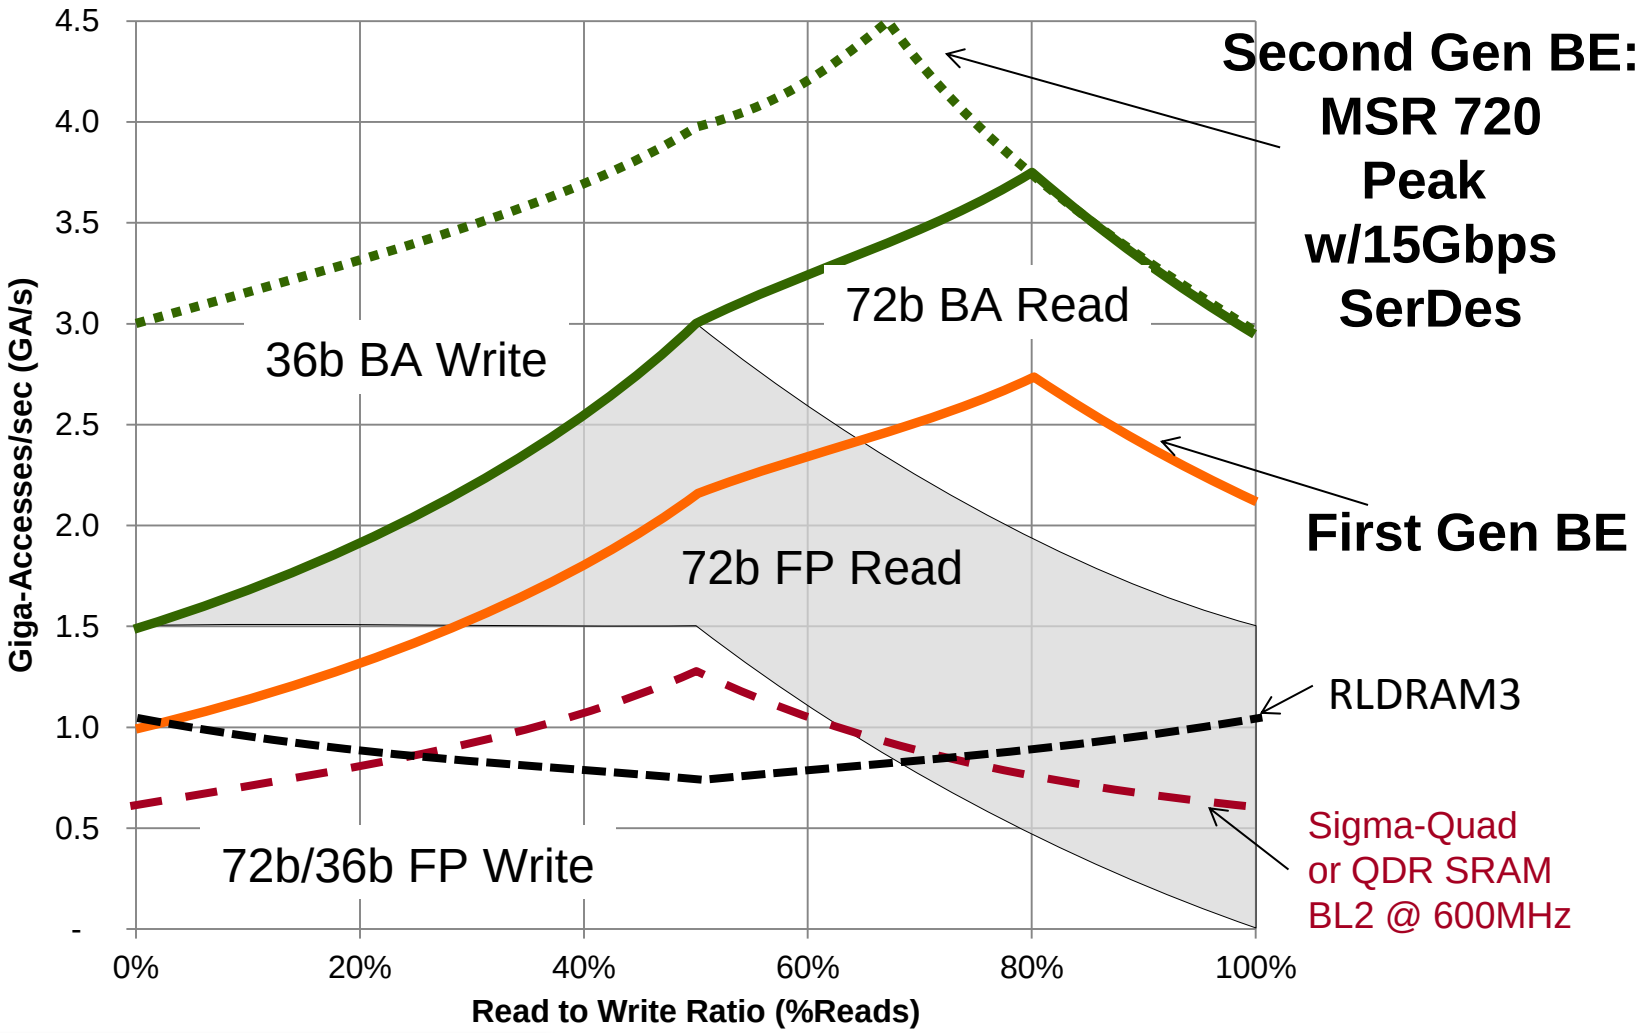

Note: Datasheet comparison - 1 access = 72 bits
- FP Access is arbitrated by BCR.
- BA Access is to/from array only. Must be bank aware.

Get More Done With The Same Amount of I/O

Bandwidth Engine 2 Architecture & Family

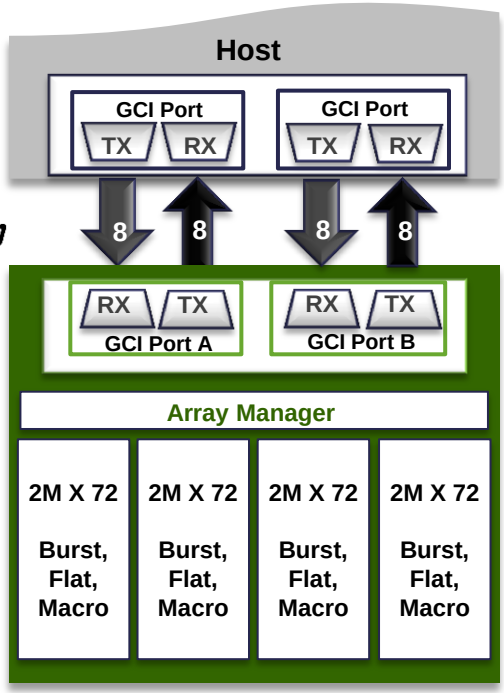

- **Parallel Array Architecture … Performance up to:**
  - 16 outstanding transactions
  - 4.5G Accesses per second
  - 192 Gbps full duplex throughput
  - ~12ns deterministic read latency
  - 2.7ns Random cycle time (tRC)
- **GigaChip Interface … 90% Efficient Transport Protocol**
  - Up to sixteen low latency SerDes lanes (8G to 15G)
- **High Reliability …70X better SER than 6T-SRAM**
  - Full ECC support; 72bit array and macro datapath operations
  - CRC protected and self recovering GigaChip Interface
  - SEU resistant 1T-SRAM Memory core: < 10 FIT/Mb
  - Bit Safe™ Self Test and Self Repair Option

		Bandwidth Engine – 2nd Generation		
Applications	**BE1 - MSR576**	**MSR620**	**MSR820**	**MSR720**
Lookup, LPM, Hash	Highest single component Access Rate			
Statistics	Onboard ALU		Onboard ALU+	
Buffer – up to 80% efficient		Per cycle Burst, Write Broadcast		
Metering, Dual Ops			Fixed Macro	
Semaphore, Link List			Atomic Ops	
State, Queuing, Link List				Dual Port w/ Data Coherency, 36 bit word access

Dual Counter: 5 Stage Pipeline

Includes Index Compare Logic (not illustrated) for Data Forwarding:

- In case of an index match, data is forwarded in the pipeline
- Prevents stale data in the pipeline
- Similar pipeline for Split Counter
- End-to-End ECC Protection
- 16 ns to completion

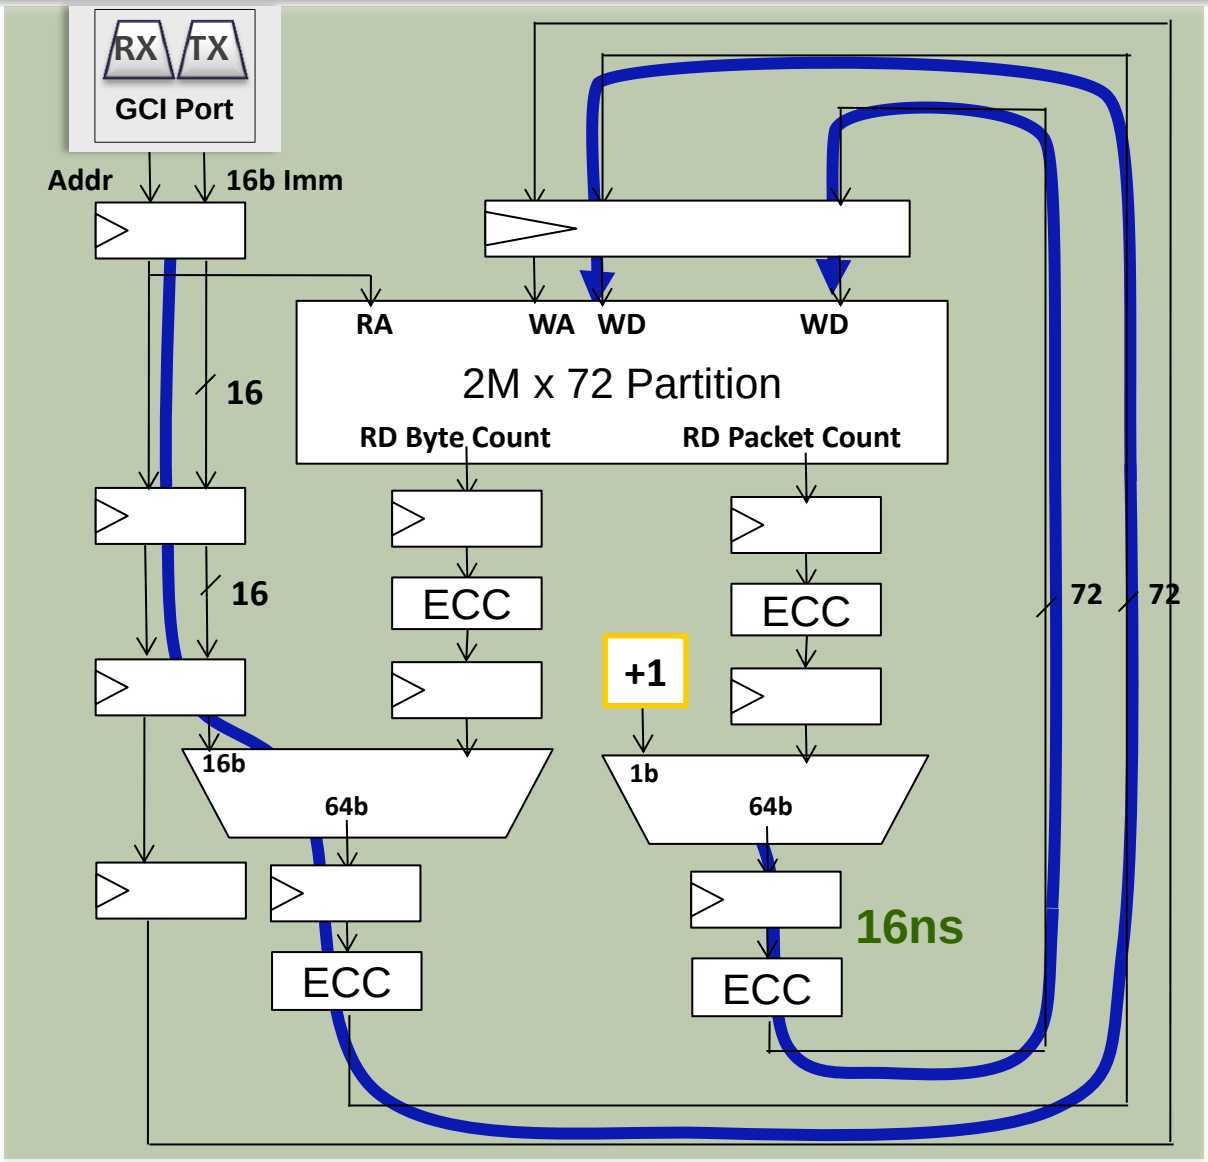

MSR820 – Metering Capabilities

- **Individual Flow Programmability**
  - Meter Type
  - Flow Rates
  - Thresholds
- **8M Two Color Flows**
- **4M Three Color Flows**
- **Line Rate 4x100G**

Basic Meter (Two Color)

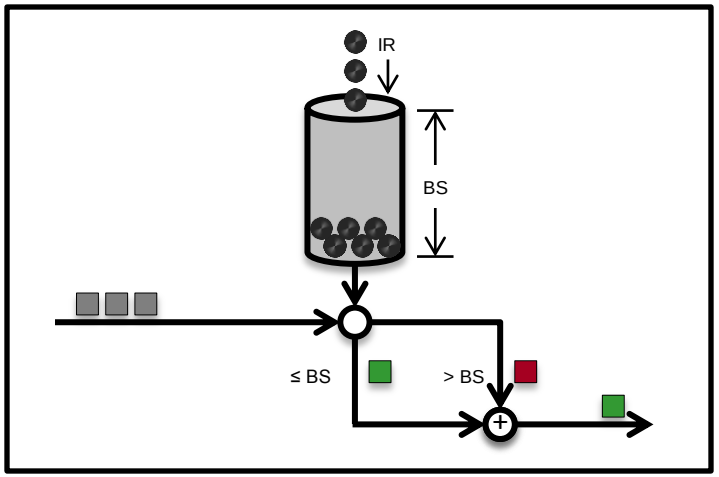

Single Rate Three Color Meter

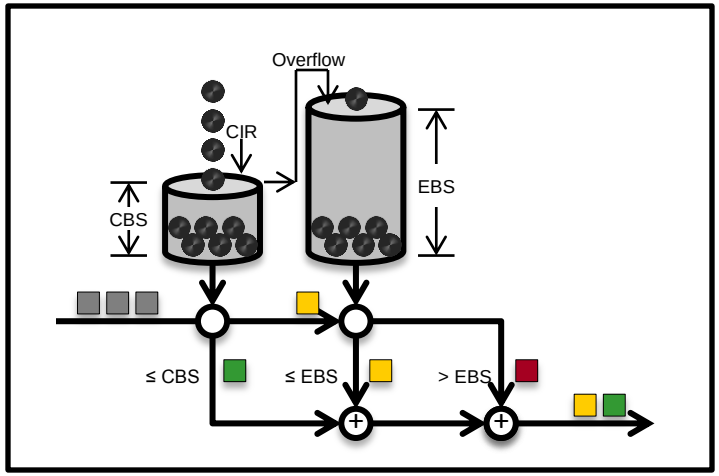

Two Rate Three Color Meter

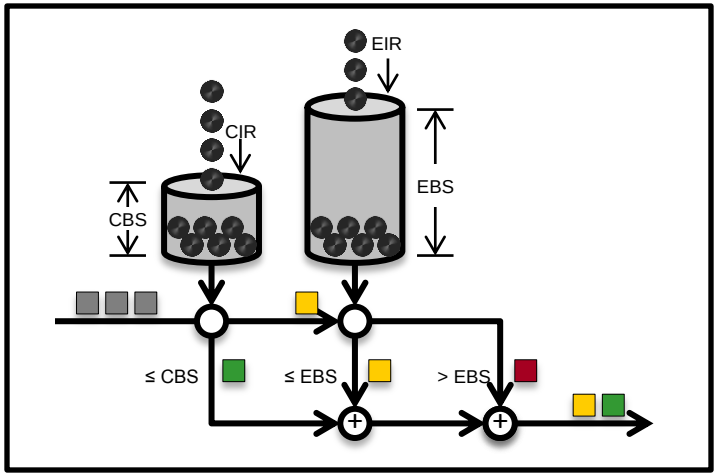

MSR820 – Bandwidth Engine 2 – Intelligent Memory Macros Leverages I/O

- **Second Generation Bandwidth Engine Architecture**
  - Up to 4.5 billion external memory accesses per second w/**16** SerDes Lanes
  - Macros support up to 6 billion internal accesses per second w/**8** SerDes Lane
  - Macros execute Atomically: Stats, Metering, Read & Set, Test & Set

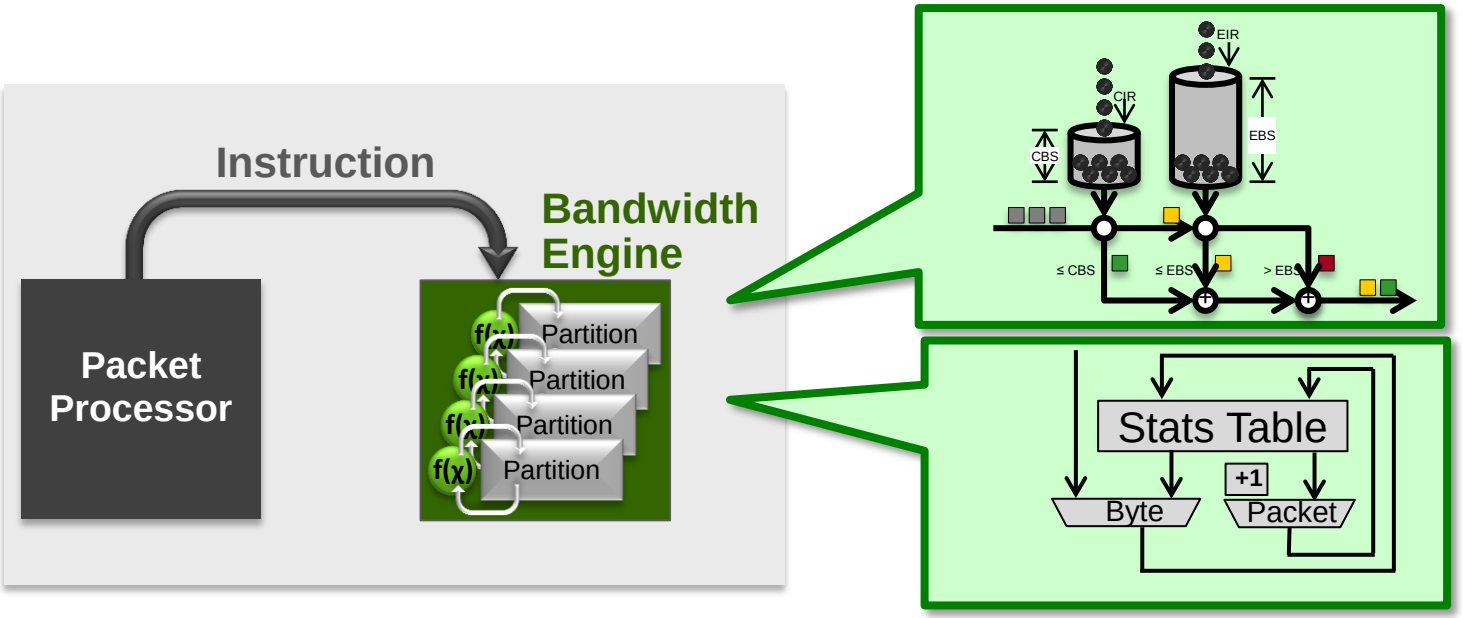

Supports:
4 x 100GE ports
w/ Stats + Metering
over 8 lanes

Intelligent Offload, Fire Forward Architecture

Physical Design

Conceptual Timing & Data Access Control

BE 1: Read Latency of 15.9ns vs. BE2: Read Latency of ~12.5ns

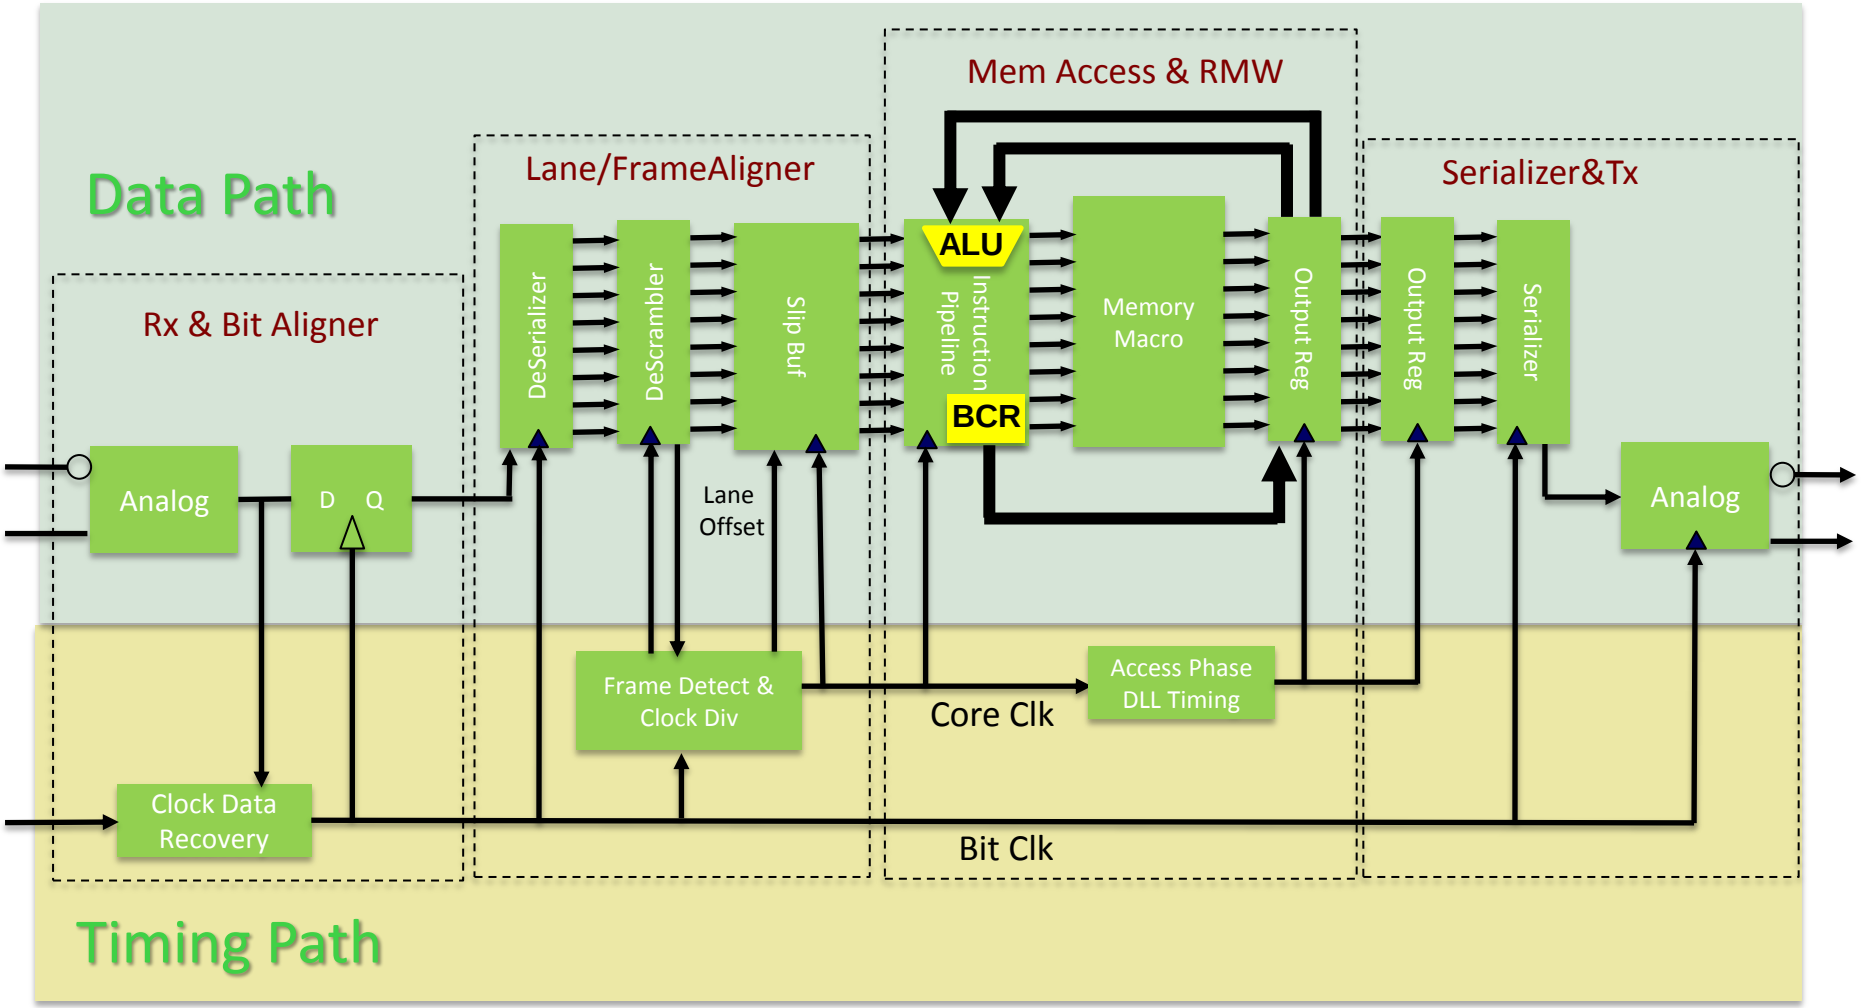

MSR720 Layout

QDR like Dual Port Performance for 1.10x vs 3x die area cost

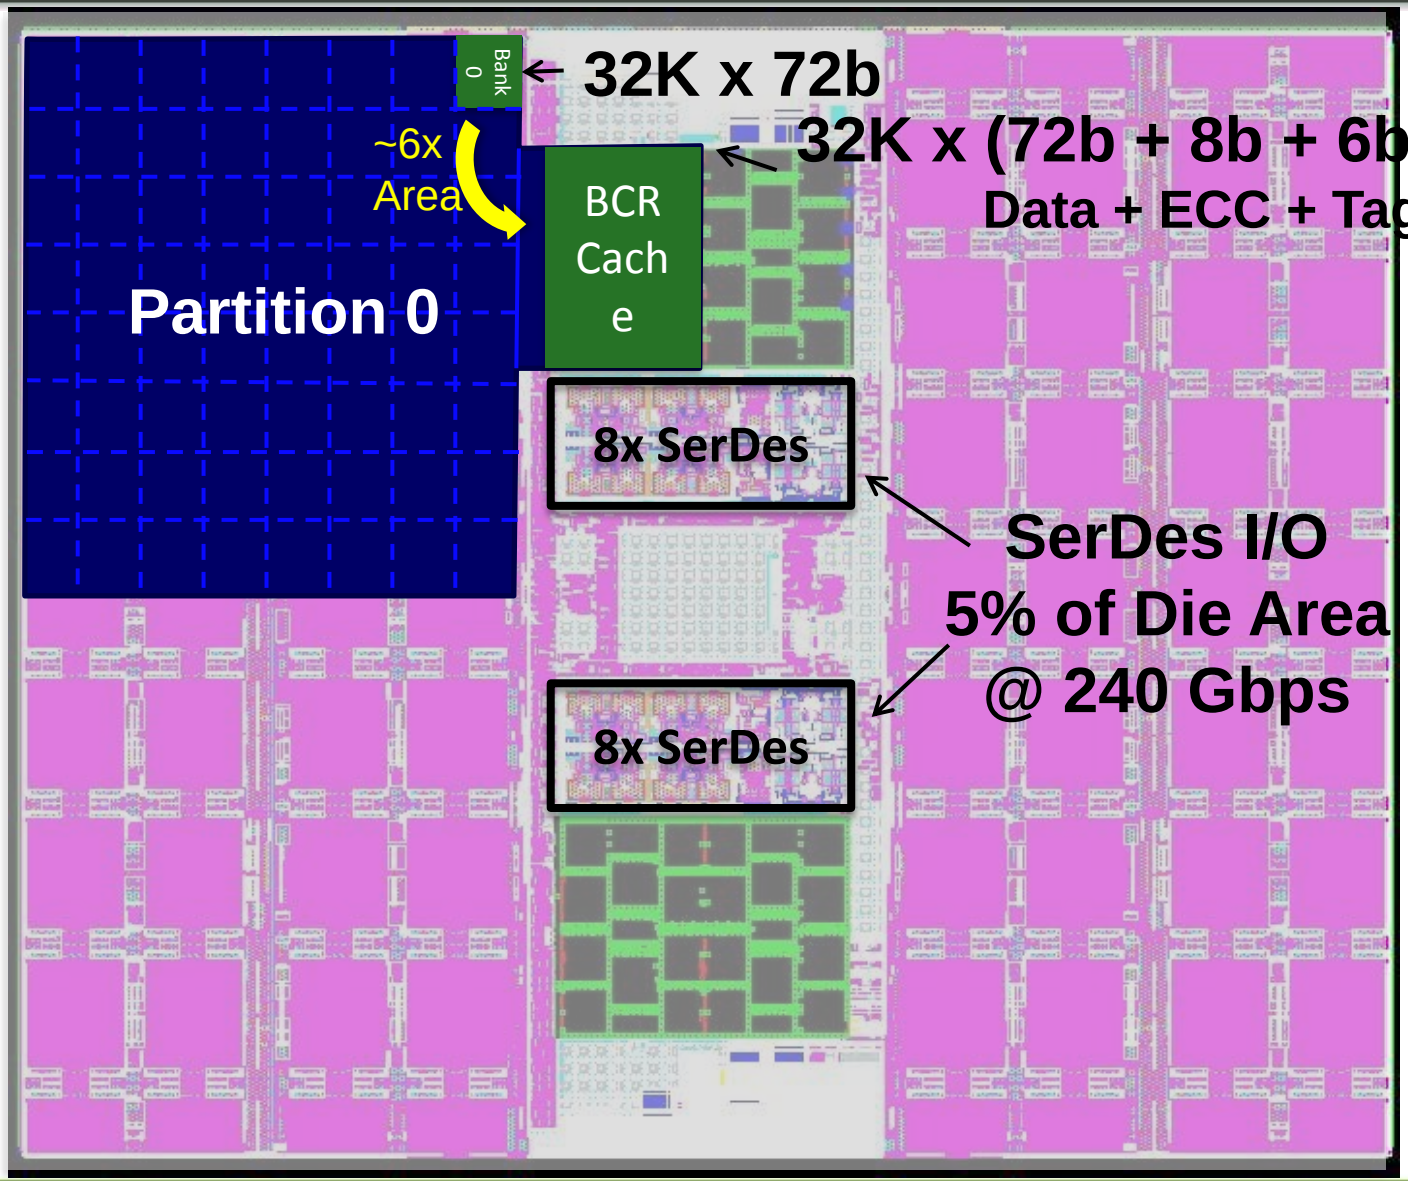

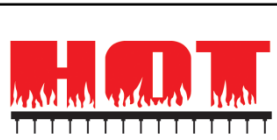

Hot Chips 25
Stanford Memorial Auditorium
August 25-27, 2013

Thank you

Michael J. Miller
mmiller@mosys.com