MoSys Accelerating Data Movement
5G Enterprise, Cloud

MoSys Legal Disclaimer

This presentation is not meant to be exhaustive and is provided as is, for convenience and information only and is not to be relied upon for any purpose, other than educational. The presentation is intended only to provide the general insights, opinions, and/or internally developed guidelines and procedures of MoSys, Inc. The information in this presentation may need to be adapted to your specific situation or work environment.

TO THE EXTENT INFORMATION IN THIS PRESENTATION IS PROVIDED IN CONNECTION WITH MOSYS PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN MOSYS'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS.

A "Mission Critical Application" is any application in which failure of the MoSys Product could result, directly or indirectly, in personal injury or death.

SHOULD YOU PURCHASE OR USE MOSYS'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD MOSYS AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS, COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT MOSYS OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE MOSYS PRODUCT OR ANY OF ITS PARTS.

MoSys assumes no liability and disclaims any express or implied warranty regarding the information in the presentation, including any liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. MoSys expressly disclaims any liability that you, or a third party might suffer due to your failure to heed the warning in this disclaimer.

MoSys makes no representation or warranties regarding this presentation's accuracy or completeness and accepts no duty to update this presentation based on more current information. MoSys retains the right to make changes at any time, without notice.

MoSys may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Any third parties referenced in this presentation are provided for information only. MOSYS DOES NOT ENDORSE ANY SPECIFIC PRODUCT OR ENTITY MENTIONED IN THIS PRESENTATION.

MoSys and the MoSys logo are trademarks of MoSys, Inc. in the U.S. and other countries. Other names and brands are the property of their respective owners.

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 (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**

Why MoSys? How Do We Solve the Pain?

We have created new ways to manage traffic flows at lower latencies

- Proven family of FPGA accelerator engine ICs
- New Virtual Accelerator Engine IP to classify packets at line rate
- New performance levels – 100s of gigabits per second
  - Can handle the millions of rules required to manage so many devices
  - Can manage the millions of routes required to route virtual flows
  - Can easily add security through millions of complex Access Control Lists
- All done in parallel, which greatly reduces latency
- Solution based on proven Intel® FPGAs – Easy plug-in solution

MoSys Stellar Packet Classification IP – What Is It?

- Ultra High-Speed Search Engine IP for FPGA
  - Includes advanced TCAM-like compiler (Ternary Content-Addressable Memory)
  - Optimized for Intel® Agilex™ FPGAs and Intel® Stratix® 10 FPGAs
- TCAM-like functions - supports n-tuple Access Control Lists (ACLs) and Longest Prefix Match (LPM)
- Up to multigigabit TCAM equivalence with millions of rules – very wide TCAMs
- Supports a wide range of Internal SRAM and external high-speed memories
- Based on MoSys Graph Memory Engine
- Very fast rule updates - atomic – no need to stop traffic

MoSys Memory & Accelerator Chip Family Overview

Communications/ Packet Processors, SmartNIC, SmartSwitch, SmartStorage…

FPGA

SerDes
Rx
Tx

QDR-like interface

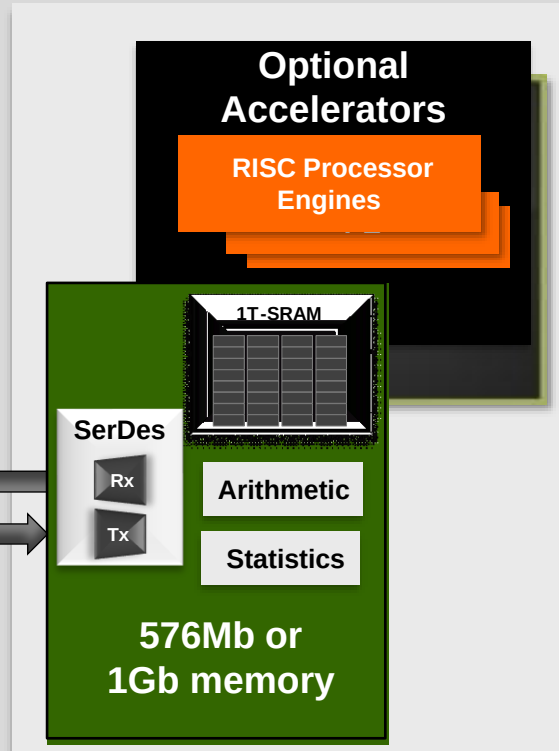

- 576Mb and 1Gb High-Speed Memory
- Extends Packet Classification Platform Capacity
- Offloads FPGA
- Very Low Latency & Power Efficient
- Low Pin Count
  - Serial Attachment but looks like QDR memory inside the FPGA
- Optional RISC engine accelerators

Adds high performance memory to FPGA for even more rule capacity

MoSys Broad Family of Memories & Accelerators

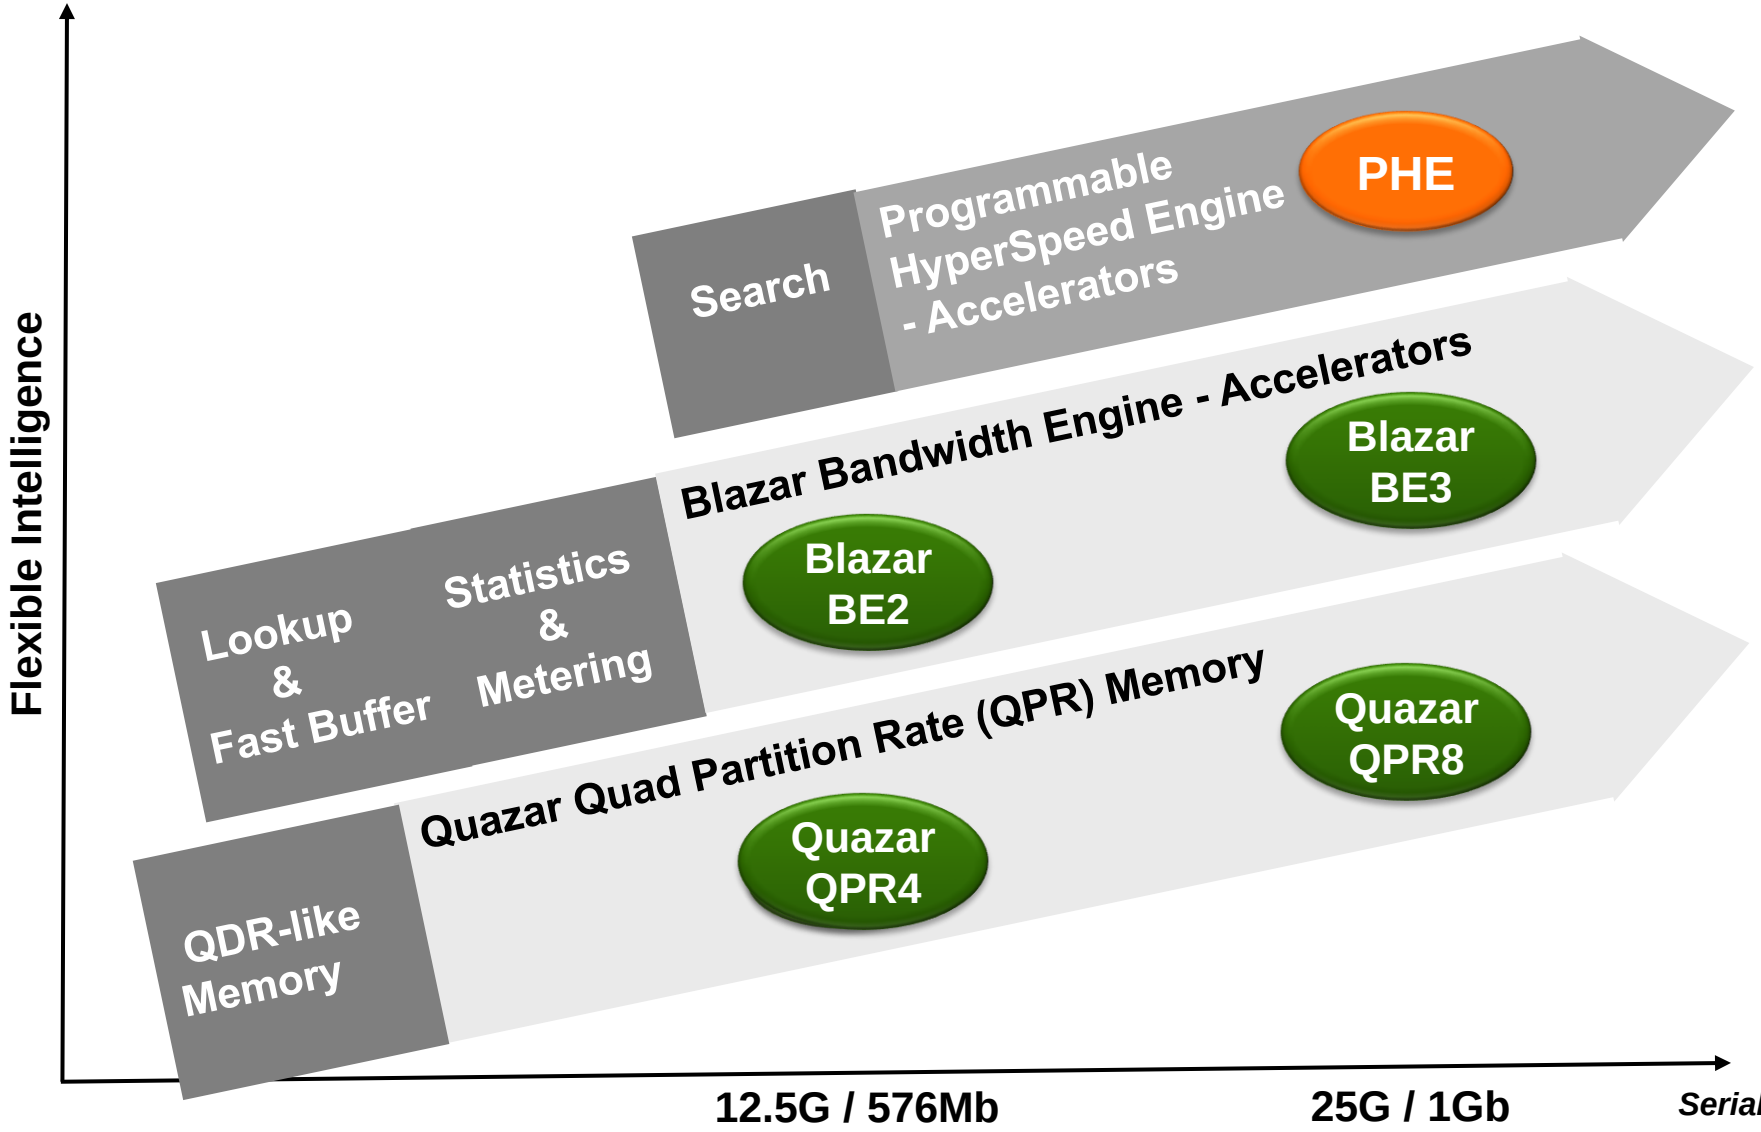

PHE includes 32 RISC cores (256 vcores)

Further Extends Search Performance & Offloads FPGA

Blazar & Quazar – Three Typical Silicon Use Cases

Delivers Highest Performance

- Highest Efficiency Interface Protocol
- Up to 480 Gbps I/O bandwidth
- 32 concurrent memory ops
- Stats include aging activity bits
- Highest Look-Up Performance
  - ➢ up to 5+ billion reads per second

1) Blazar Adds Stats & Metering – *Up to 32 x 100G*

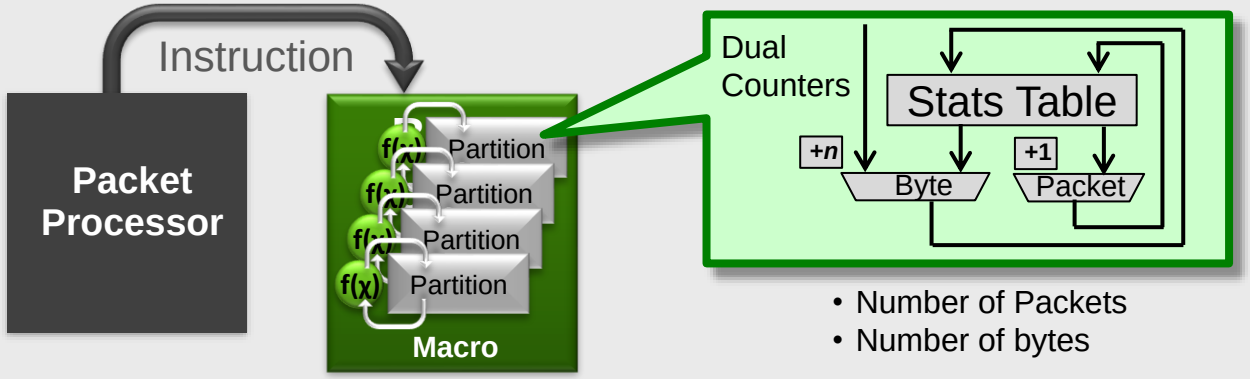

2) Blazar & Quazar Buffer - *Up to 400G Full Duplex*

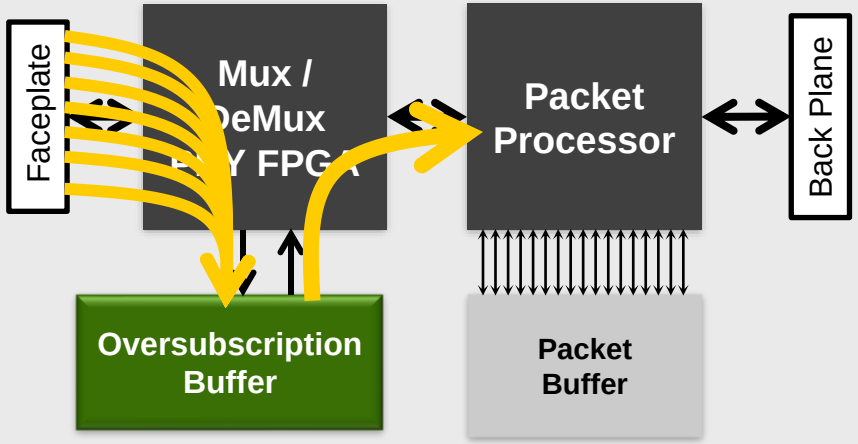

3) Blazar & Quazar Packet Classification Platform - Table Lookup *Up to 5B Read/s*

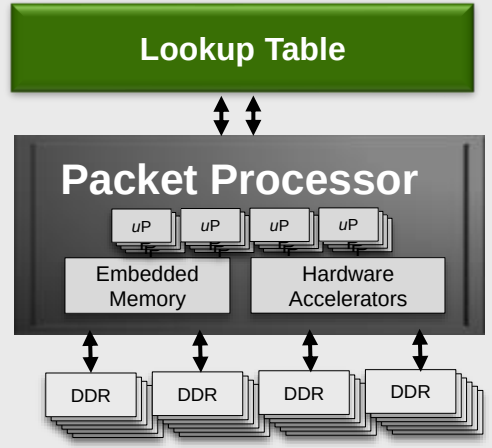

1st Polling Question

- **Q1) Do you have high bandwidth FPGA applications (80Gbps or above aggregate line rates)?**
  - I design 80Gbps or above aggregate line rate boards
  - I am considering using TCAMs
  - My application includes Packet Header Inspection, Payload inspection, Look Up Tables
  - My application may need small (5-10ms) ingress or egress smoothing buffers

Accelerator Engine Performance Based on Architecture

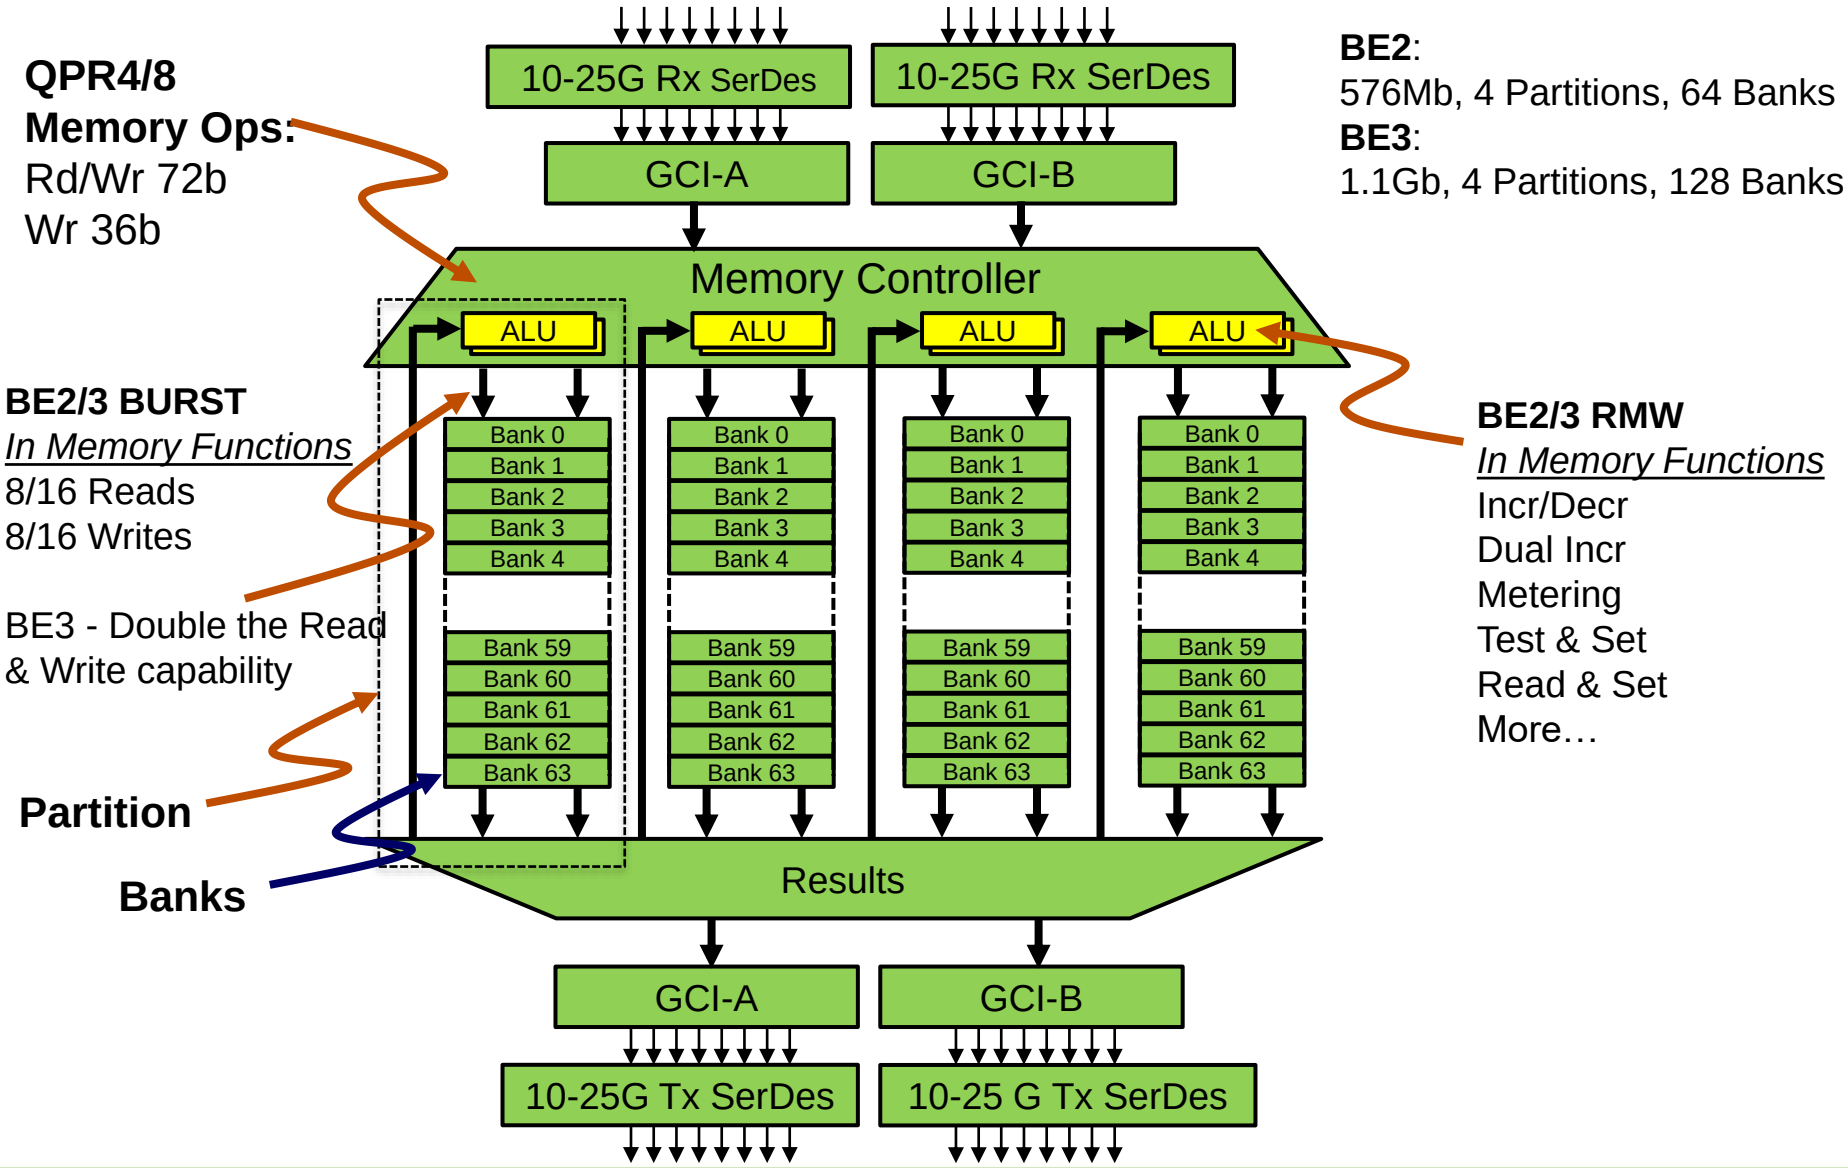

BE2:
576Mb, 4 Partitions, 64 Banks
BE3:
1.1Gb, 4 Partitions, 128 Banks

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)

Pipelining With BE3 As A Dual Port

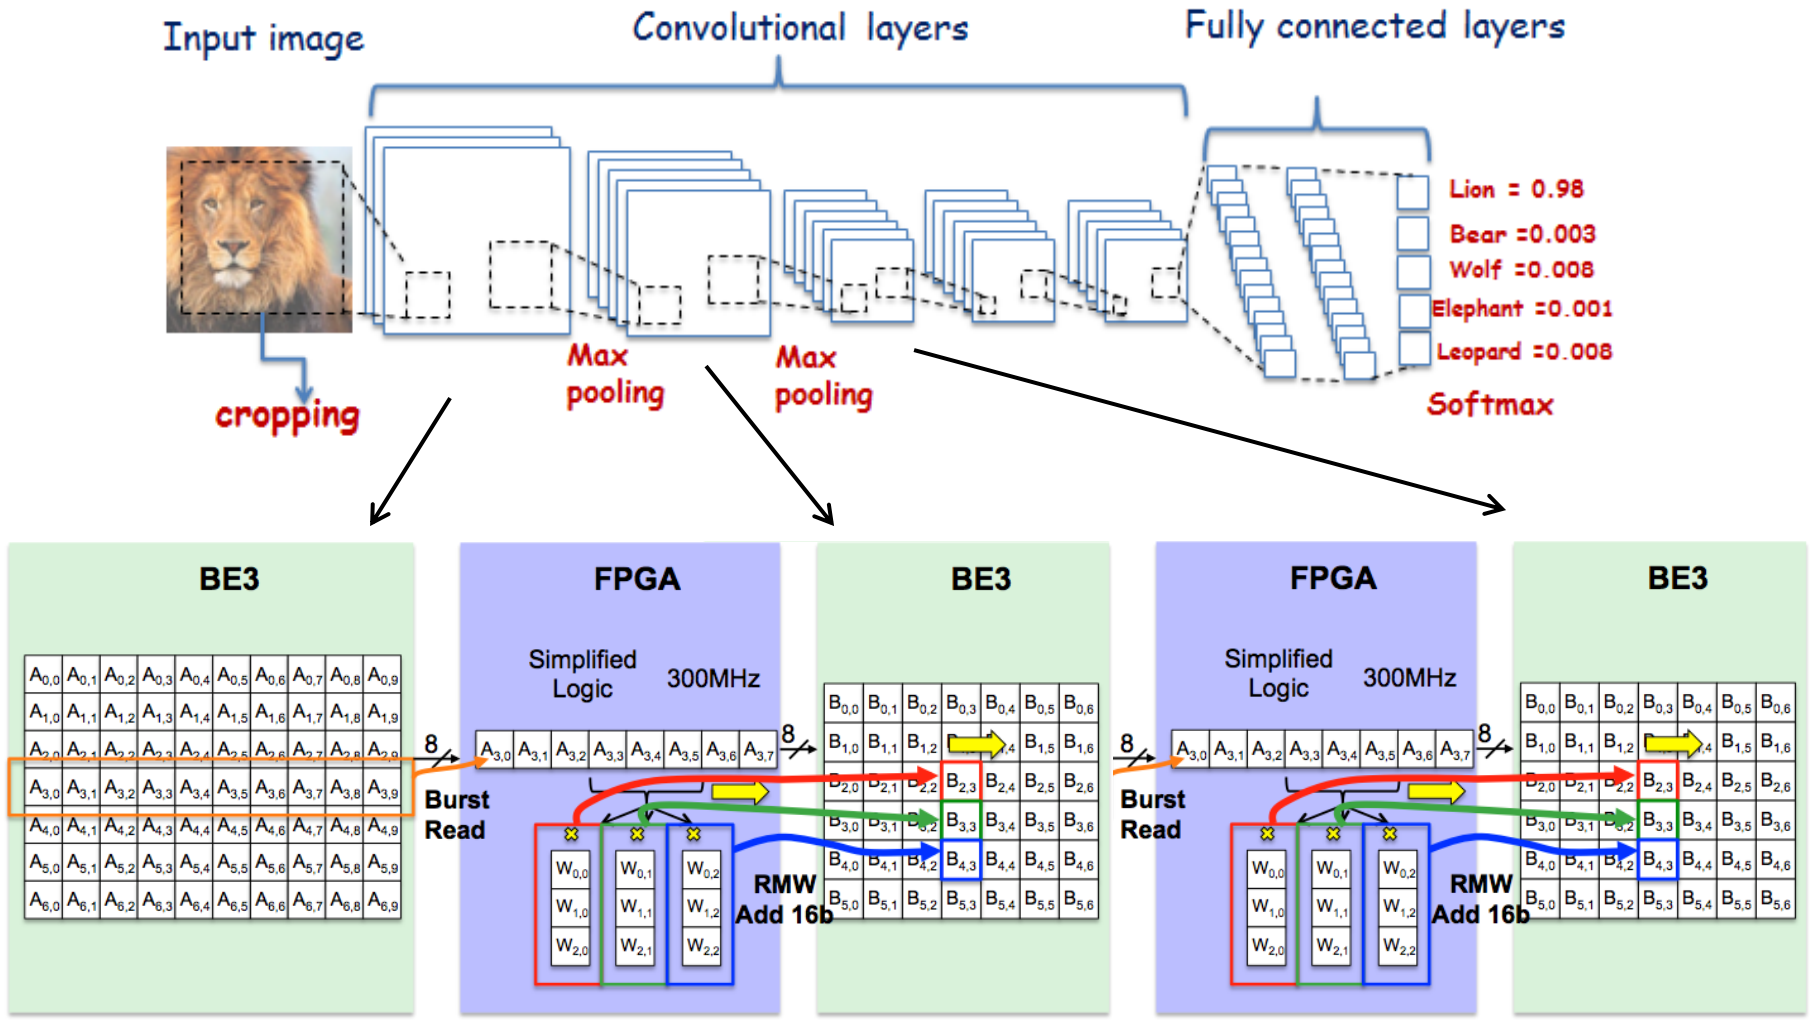

Memory Architecture is Critical
Increase Performance AND Save Space/Cost

Parallel vs Serial

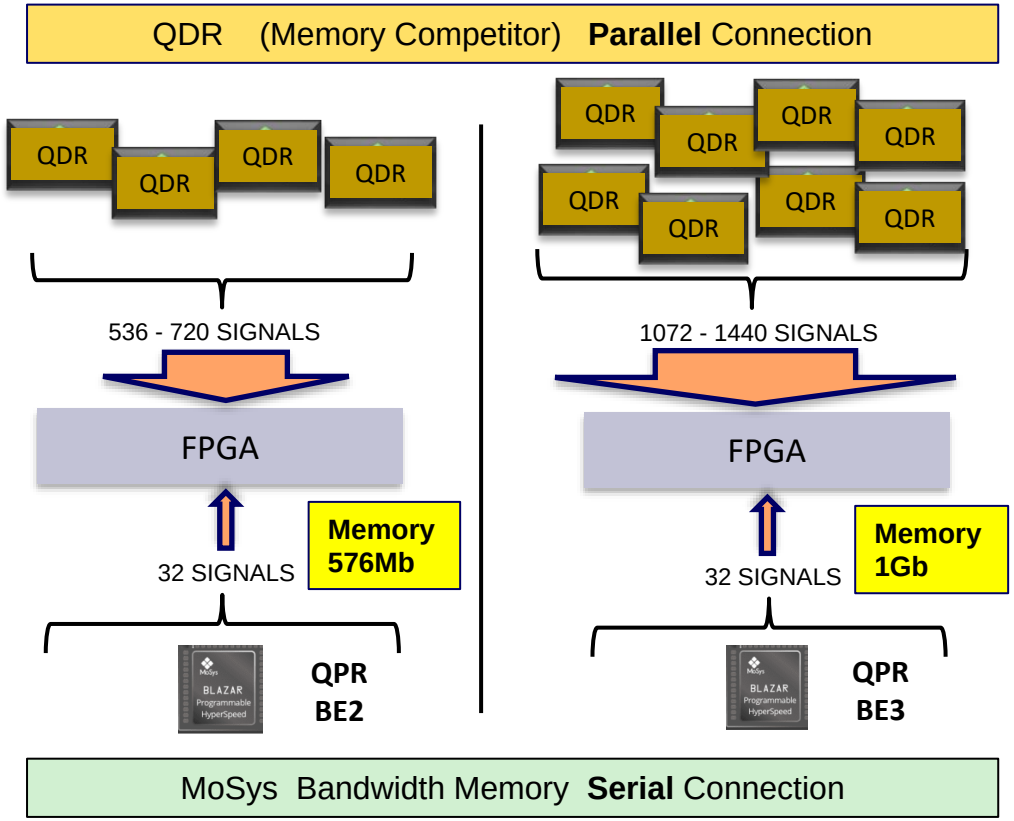

Space & Cost Considerations

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 576Mb 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

3rd Gen Memory Architecture
16 Dual Counter ALUs ➔ 5B RMW Ops/s

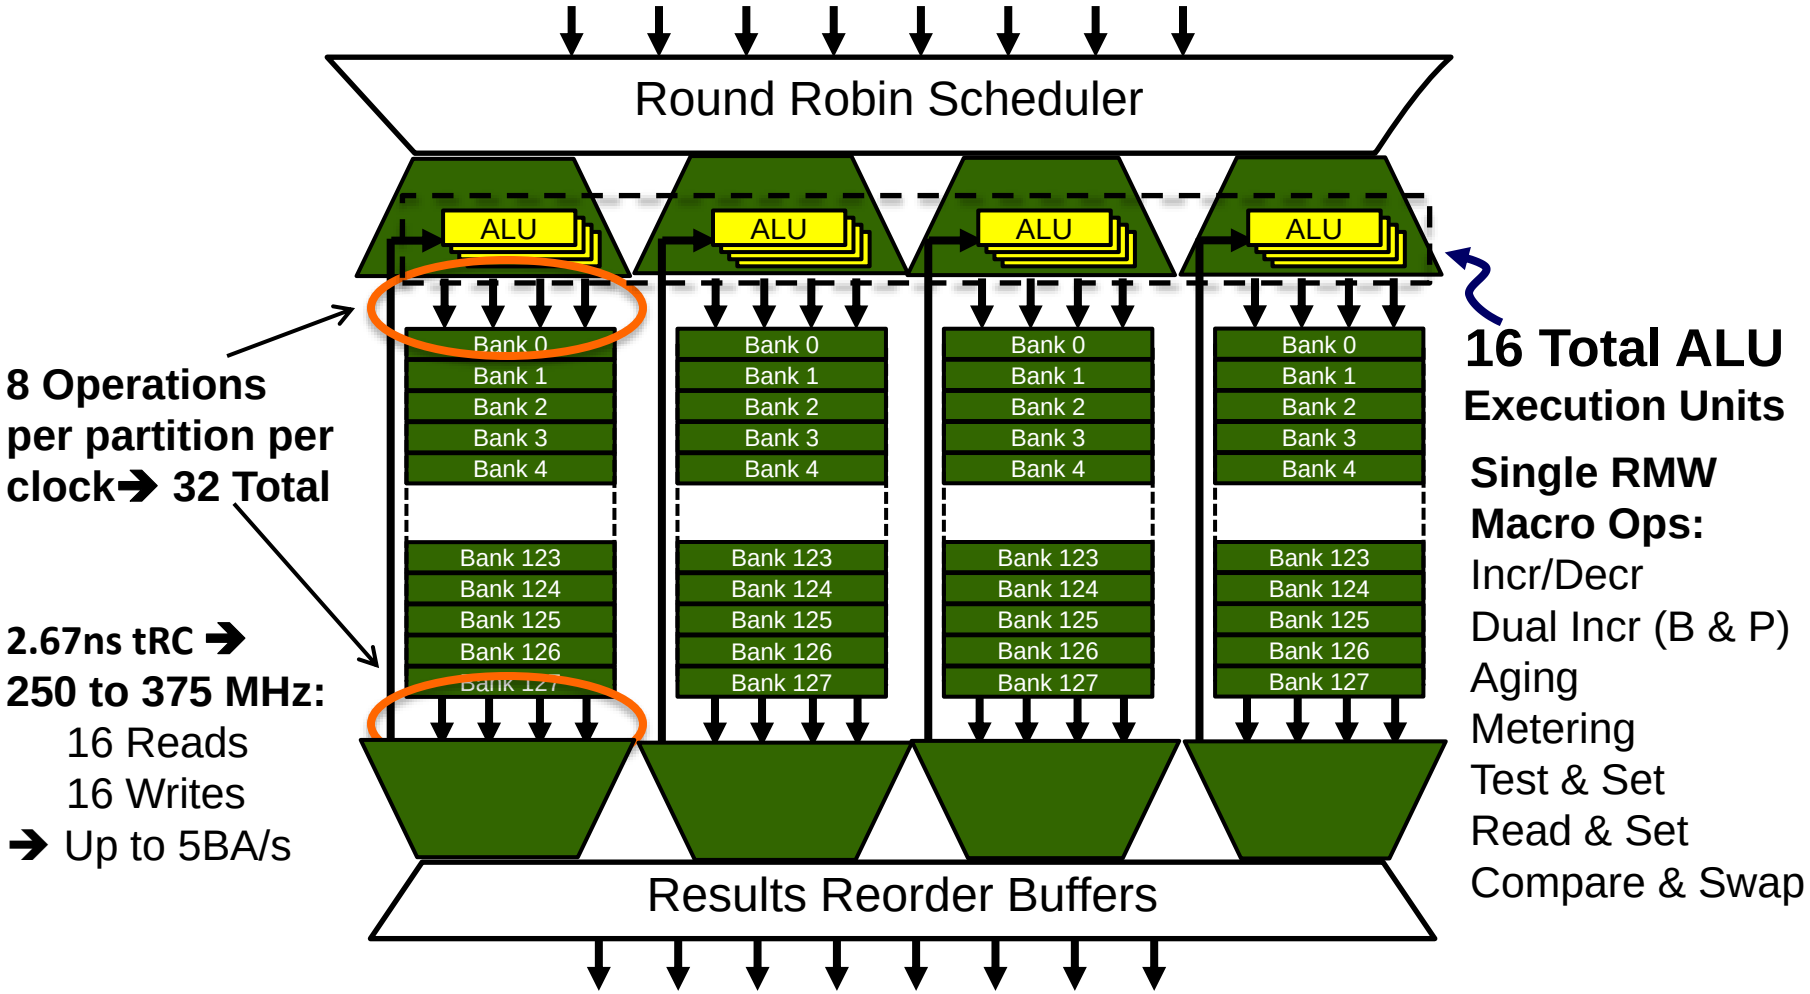

Programmable HyperSpeed Engine (PHE) Architecture

Physical
- 16 x 10 to 30Gbps 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

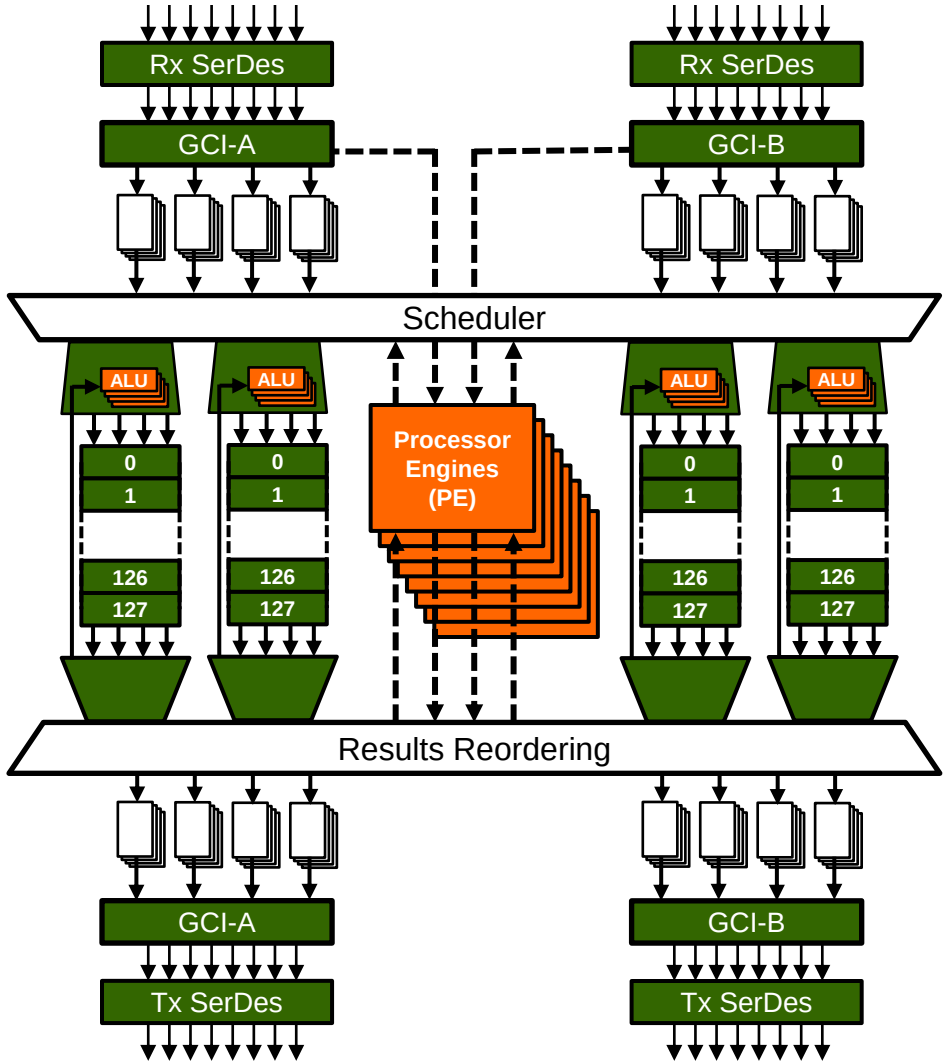

Instruction Set Overview

- **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

Summary

- **Bandwidth Engine 2 and 3 in full production**
- **Over 200K shipped, no field failures**
- **Qualified as Enterprise Grade and ships to some of the largest networking and security companies in the world**
- **Replaces between 4 and 8 QDR devices**
- **Price savings**
- **Board design time savings**
- **Proven interop with Intel and Xilinx FPGA devices**
- **Eval boards available**
- **MoSys supports custom controller development for customers**
- **Available in Commercial and Extended temp**
  - 0C to 85C
  - 0C to 100C
- **Bandwidth Engine 3 is available in Industrial Temp**
  - -40 to 85C

2nd Polling Question

Q2) Are you working on a 5G design ?

- Yes in the next 6 months using a custom design
- Yes in the next 6 months designing to a Service Provider spec
- Yes in the next 12 months using a custom design
- Yes in the next 12 months designing to a Service Provider spec
- No not working on a 5G specific solution

Accelerating Data Intelligence

MoSys Stellar
Packet Classification Platform
Overview – Non-NDA

John Bromhead
VP of Bus Dev
jbromhead@mosys.com
+1-858-642-2501

Stellar Packet Classifier Platform Exec Summary

- **Ultra High-Speed Search Engine IP** for ASIC/FPGA plus advanced TCAM-like compiler
- **TCAM-like functions** - supports n-tuple ACLs & LPM
- **Can be optimized to work with a wide range of High-Speed Memories**
  - ASIC SRAM, FPGA M20K, eSRAM, BRAM, URAM, DDR4, HBM and/or other external memory
- **Graph Memory Engine** – Allows easy mapping of complex rules
- **Very fast updates** - can be dynamically updated on the fly – incremental or full updates
- **Flexible** – ASIC/FPGA only or ASIC/FPGA plus external memory i.e. MoSys memory or HBM
- **Performance**
  - **Up to multigigabit TCAM** equivalence with **millions of rules**
    - can greatly extend internal Switch and SmartNIC TCAM functions
  - **Hundreds of Millions searches per second**
  - **Hundred of Gigabits per second**

210714

Why MoSys? How Do We Solve the Pain?

We have created new ways to manage traffic flows at lower latencies

- Proven family of FPGA accelerator engine ICs
- New Virtual Accelerator Engine IP to classify packets at line rate
- New performance levels – 100s of gigabits per second
  - Can handle the millions of rules required to manage so many devices
  - Can manage the millions of routes required to route virtual flows
  - Can easily add security through millions of complex Access Control Lists
- All done in parallel, which greatly reduces latency
- Solution based on proven Intel® FPGAs – Easy plug-in solution

5G Core & Edge Use Cases – UPF Locations

- UPF appears in 2 places
  1) Core Network
  2) At the Edge inside a MEC (Multi-Access Edge Controller)
- Often based largely on SmartNIC or SmartSwitch as software only cannot handle required performance
- Control info comes for SMF using Packet Forwarding Control Protocol (PFCP)
- SMF acts as a sort of SDN Controller
- UPF can contain **millions** of rules
  - Non learned
  - All come from SMF

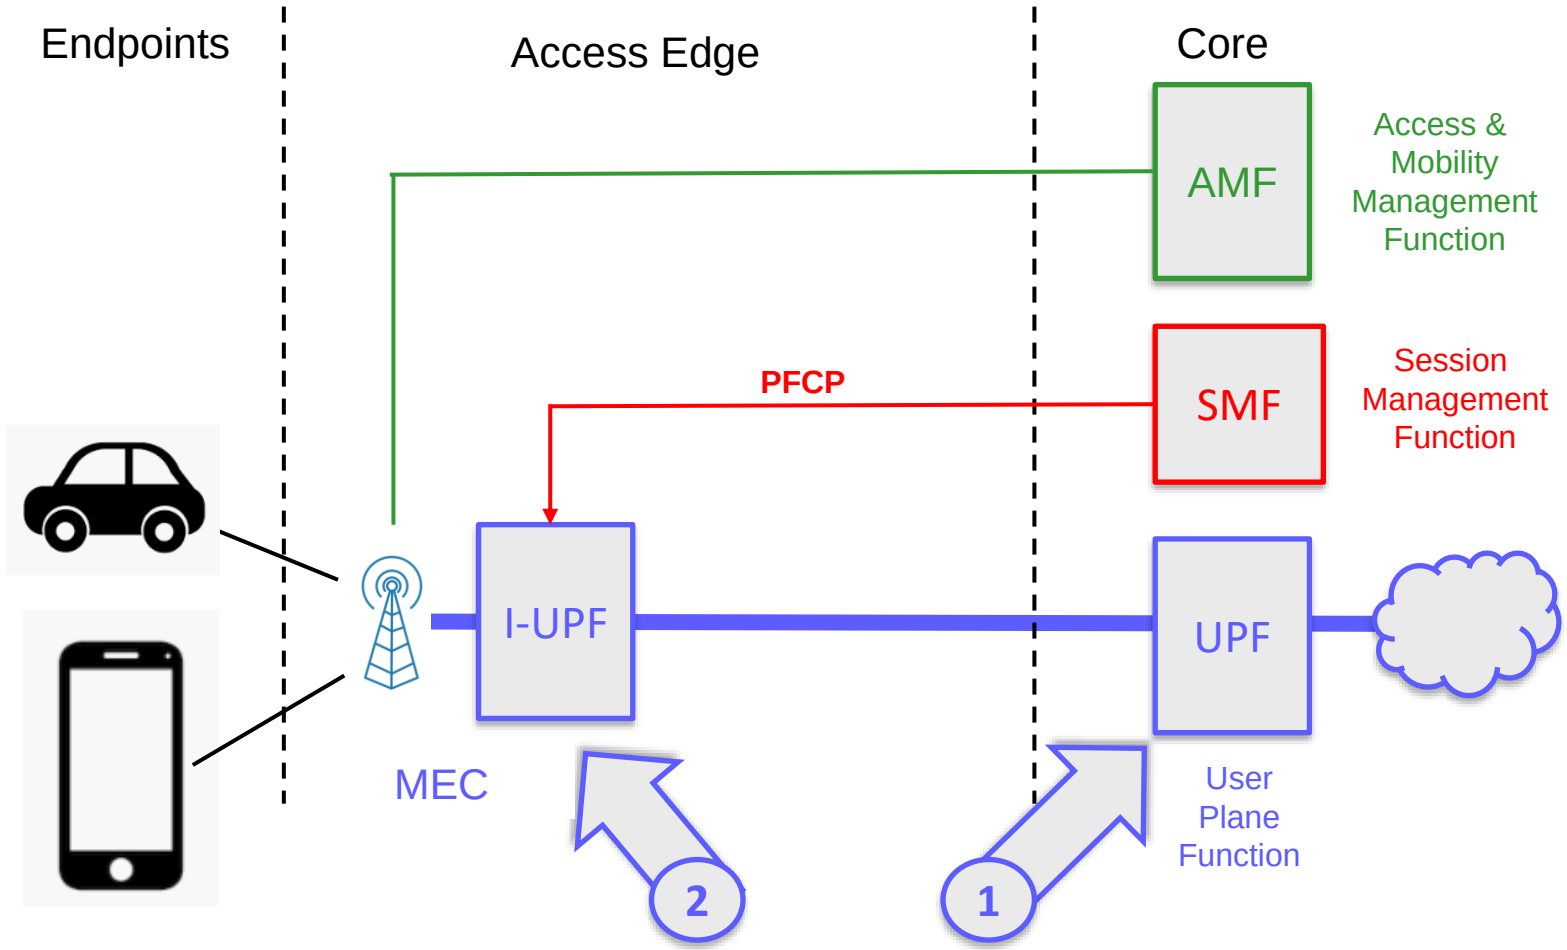

Simplified diagram (not all connections shown)

Packet Processing Architecture MoSys Stellar IP

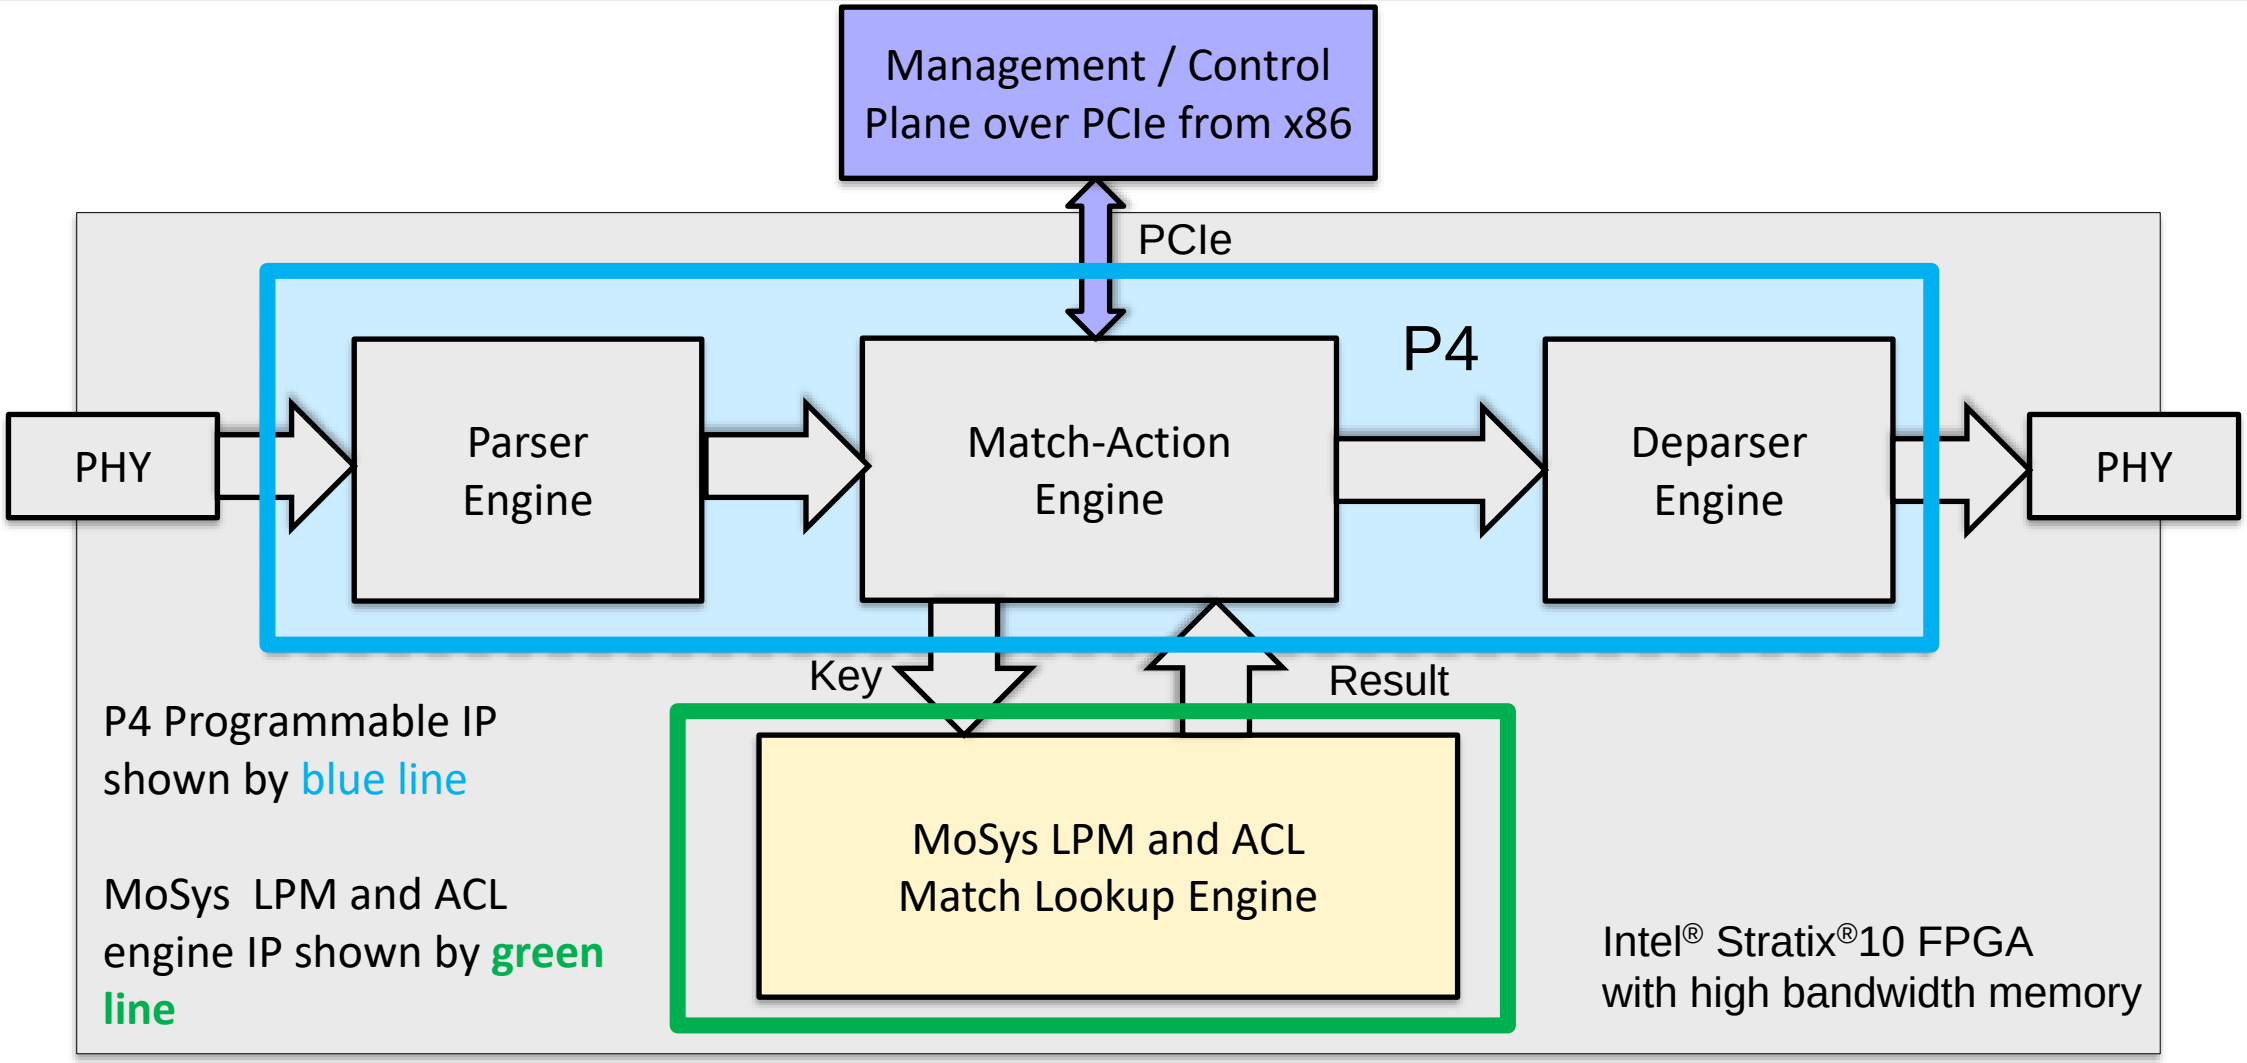

MoSys Stellar – IP Product Lines

MoSys Stellar Packet Classification Platform
Two main IP product lines both take advantage of our Scalable Graph Memory Engines (GME)

1) High Flexibility / High Complexity ACL & LPM IP

- Up to 200 Gbps at 300 MHz FPGA*
- 1 – 10+ tuple, 40 - 480b typical key size
- ACL rules up to 16+ million
- Up to 150 million lookups per second

2) High Performance / High Capacity LPM IP

- Up to 400 Gbps at 300 MHz FPGA*
- 1 – 2 tuple, 40 - 160b typical key size
- LPM rules up to 4+ million
- Up to 300 million lookups per second

Choose 1 for higher complexity
Choose 2 for higher performance

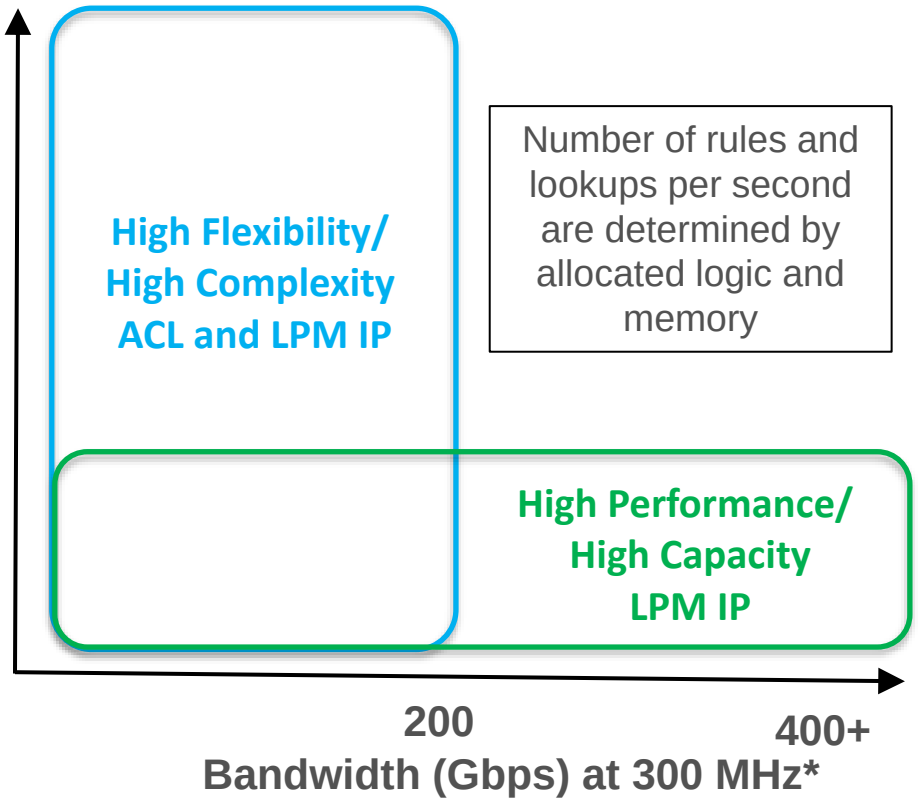

GME Scalable Performance

- Performance is scalable to meet your rule and lookups/sec requirements
- Multiple Graph Memory Engines (GME) can be combined for other configurations
- Parameterizable to optimize for bandwidth, lookup rates, number of rules or logic/memory usage

Licensing Model – One time license plus per use royalty, annual maintenance

*up to 3X with ASIC at 900 MHz

Subject to change

3rd Polling Question

- **Q3) In your 5G application do you need Packet Classification IP for your FPGA ?**
  - I would prefer to use FPGA RTL IP and on chip SRAM and/or external DRAM to do Packet Classification and Search vs. external TCAMs
  - I will implement Exact Match, LPM and/or ACLs
  - I need to use HBM already so Packet Classification IP that can leverage this would be valuable
  - I will implement Security functions like Firewalls, Anti-DDoS

Stellar Scalable Performance – 1x to 100x+ Range

- **MoSys Packet Classification Platform can run on variety of Graph Memory Engines (GME)**
  Available as Software, FPFA/ASIC RTL or Firmware for PHE
- **Scalable across many platforms**
  - From "C" → ASIC/FPGA IP core + Internal SRAM and optional DDR/HBM/MoSys BE2/BE3/PHE memories → RISC cores
  - Same high-level software interface across all platforms
  - Memory allocation is transparent to SW - eases porting

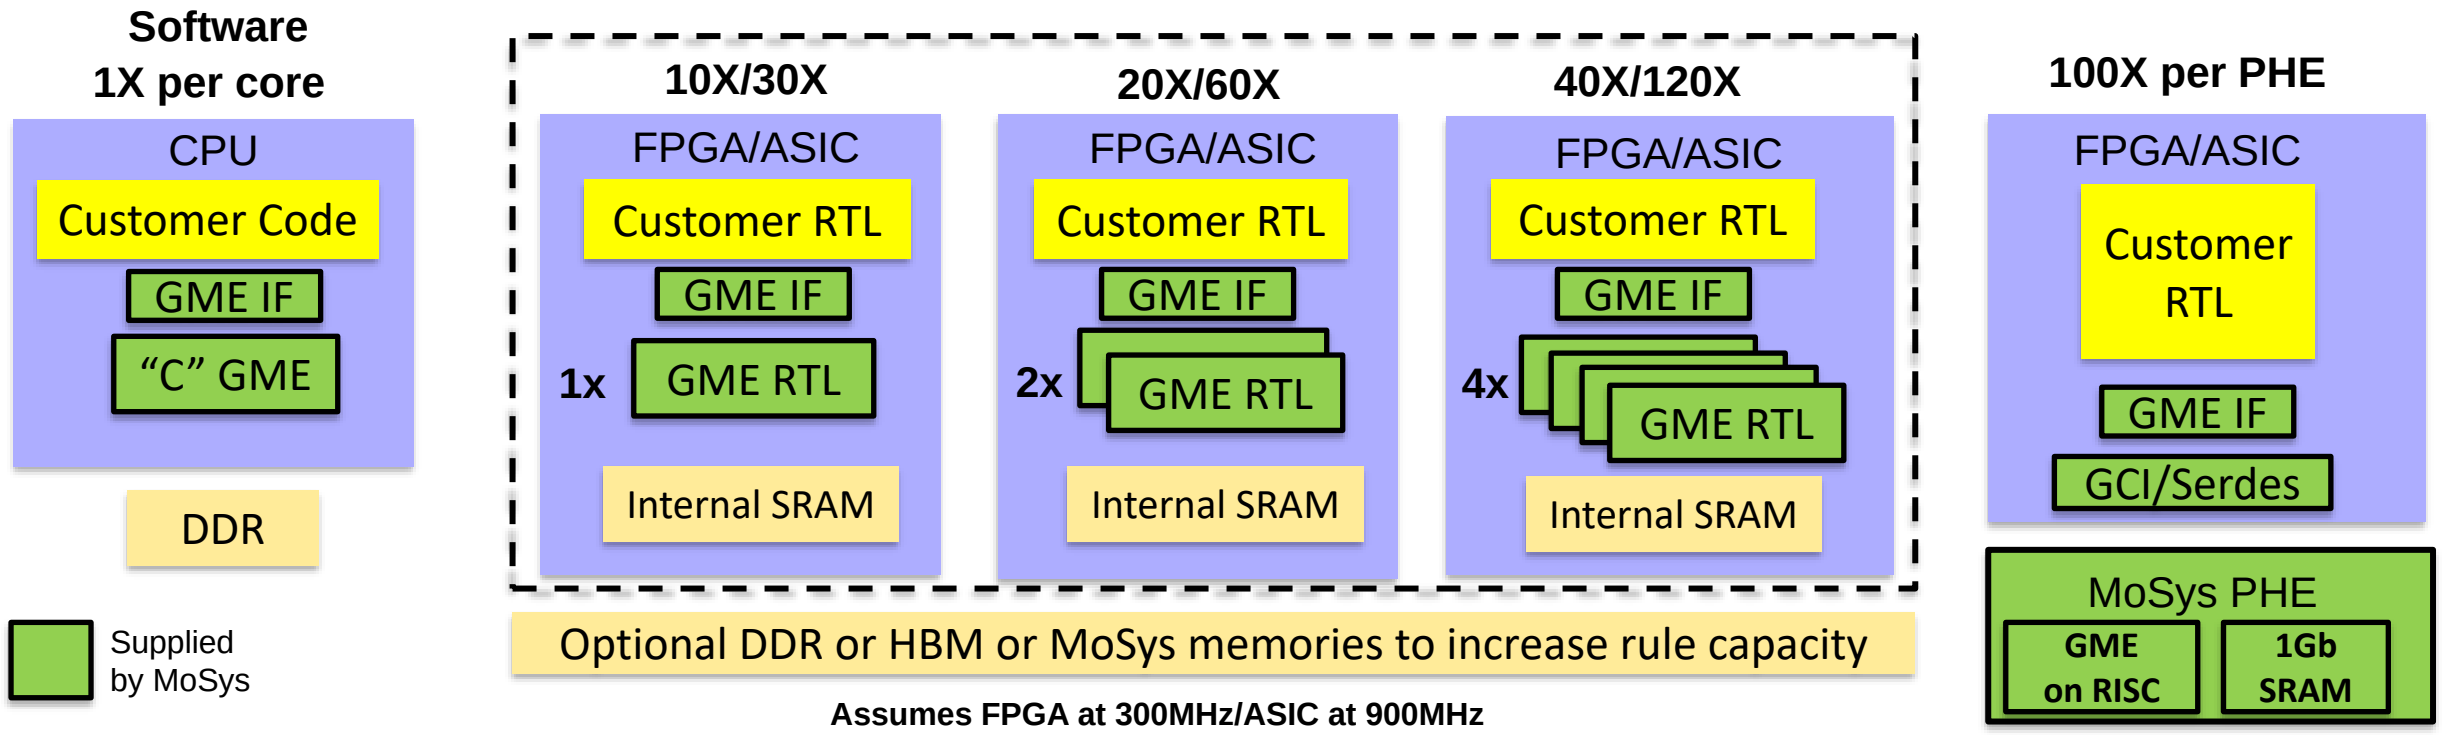

MoSys Graph Memory Engine (GME)

- **GME uses graph to inspect packet header**

```
#     0 [******** ******** ******** ********
#     1 [******** ******** ******** ********
#     2 [******** ******** ******** ********
#     3 [10111011 00000011 00010111 ********
#     4 [10111011 00000011 00010010 ********
```

Sample set of TCAM style rules
1, 0 and * (don't care)

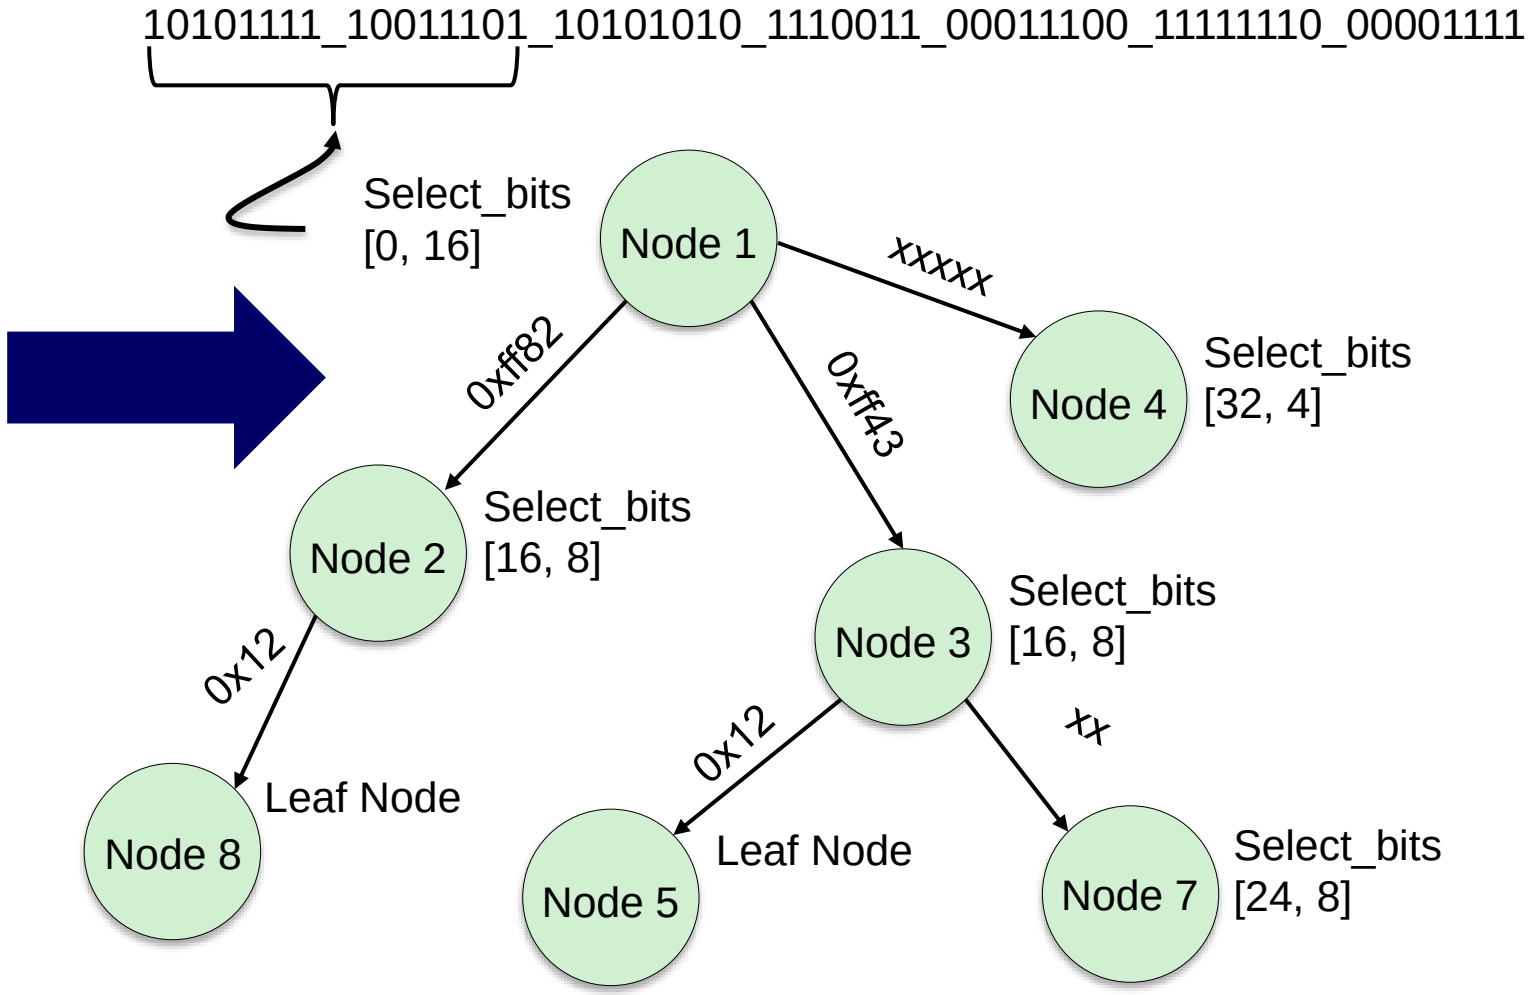

MoSys Multi-Field Multi Match Example

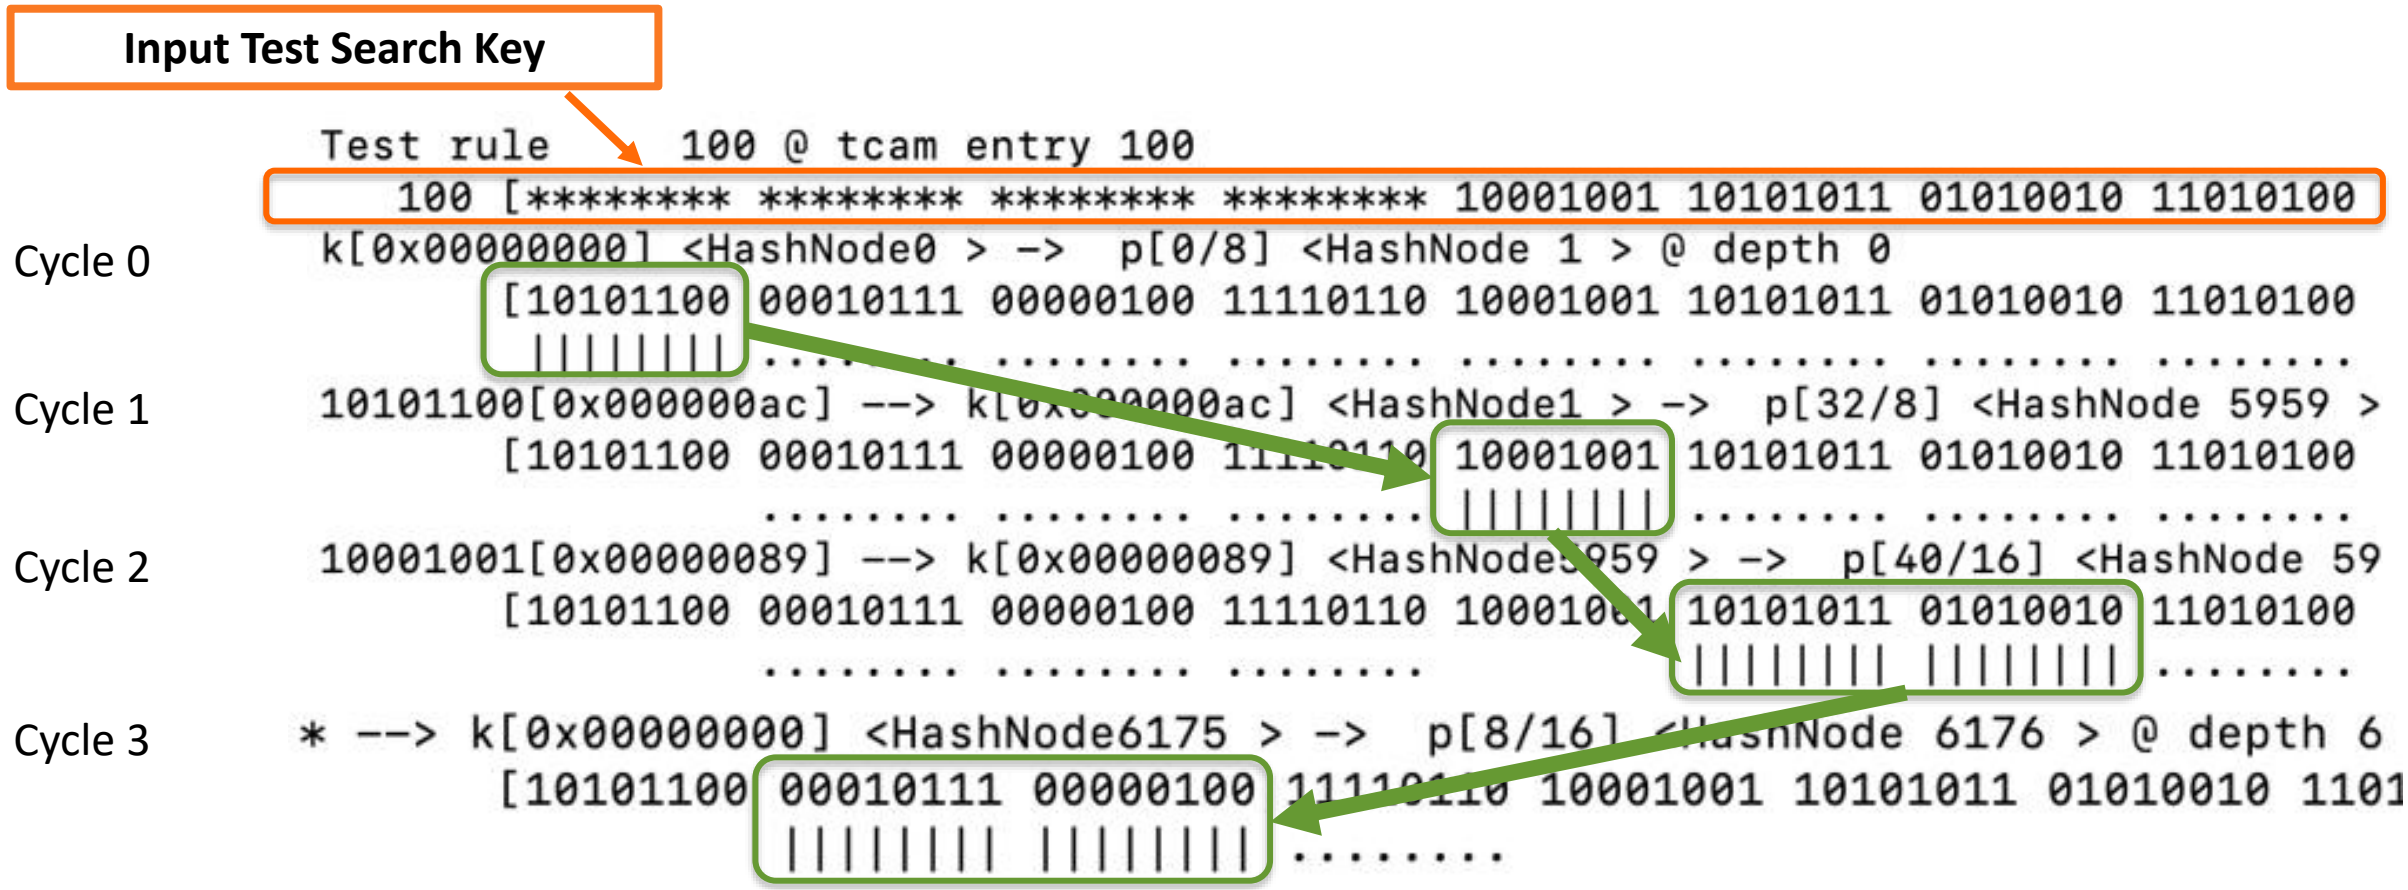

Can have multiple rules that match

Stellar IP sorts out which is the best match based on rule priority

Thank You

Gus Lignos
V.P. Sales
glignos@mosys.com

Mark Baumann
Dir of Applications
mbaumann@mosys.com

Backup – LineSpeed PHY ICs & Accelerator Engine IC part offering guide

LineSpeed™ 100G PHY IC Family

Flexible Line Card & Module Solutions

LineSpeed™ 100G PHY ICs

Applications

- Extending Reach of High-Speed Serial Links
  - Retimers with up to 20 lanes at 28Gb/s
- QSFP28-based Optical Interfaces
  - 100G (4x25G) Retimer w/ and w/o RS-FEC
  - 100G Gearbox w/ and w/o RS-FEC
- Bringing High-Speed Links to low-cost FPGAs
  - Mux/Demux of low-speed I/Os into high-speed I/Os
- Breakout of 100G ports to 10 x 10Gb Ethernet
  - Multi-Link Gearbox (MLG), including SyncE support
- High-Availability Systems with Redundant Links

Key Features

- Family of Retimer, Gearbox, & Mux/Demux ICs
- Compliant to IEEE, ITU, and OIF standards
- 100G RS-FEC encoder/decoder (optional)
- Redundant Link Mode for Protection Switching
- Independent baud rates per lane
- Strong, self-adapting receive equalizers
- Built-in PRBS generation and checking
- Multiple packages for line cards and modules
- *Priced at less than $50 in volume*

Product Description					Functions								Reach	Supported Rates	
	Part Number	Description	TX/RX Lanes	Package	Gearbox	MLG	Retimer	10x10G Retimer	4x25G Retimer	Clause 91 RSFEC	Mux/ Demux	Redundant Link Mode	15-20dB w/o FEC	10-14G	25-28G
Retimers	MSH221SF	100G Octal Retimer w/ FEC	8	12x12mm			✓		✓	✓			✓	✓	✓
	MSH222S	100G Full Duplex Retimer	8	13x13mm			✓		✓				✓	✓	✓
	MSH222SF	100G Full Duplex Retimer w/ FEC	8	13x13mm			✓		✓	✓			✓	✓	✓
	MSH225S	10 Lane Full Duplex Retimer	20	17x17mm			✓	✓					✓	✓	✓
Gearbox	MSH320S	100G Gearbox	20	17x17mm	✓			✓					✓	✓	✓
	MSH320SF	100G Gearbox w/ FEC	20	17x17mm	✓			✓		✓			✓	✓	✓
	MSH321S	100G MLG Gearbox	14	12x12mm	✓	✓			✓				✓	✓	✓
	MSH322S	100G MLG Gearbox	14	17x17mm	✓	✓			✓				✓	✓	✓
Mux	MSH420S	10:5 Mux/Demux	20	17x17mm							✓	✓	✓	✓	✓
	MSH422S	4:2 Mux/Demux	8	13x13mm							✓	✓	✓	✓	✓

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