

# A New Memory Channel

A Microprocessor Standards Committee  
Study Group

Bob Davis

bob@scsi.com

Summit Computer Systems, Inc.

January 19 2004

# Agenda

- Why Design a New Memory Interface?
- What's Changing
- Memory Hierarchy
  - Regularizing the Stored Data Interface
- Memory Vendor Design Freedom
- An Possible “Ideal” Interface
- Memory Connection Capabilities
- Possible Memory Interface Solutions
- Who Benefits from this approach?
- Issues

# Why a New Memory Interface?

- Support the Changing Memory Needs
- Extend the Addressing Capability
- Improve Reliability (RAS) – Support Redundancy
- Allow design freedom for Memory Vendors
- Recognize Memory Hierarchy and use
- Regularize storage interconnect
- Improve Signal Integrity and usage models
- Reduce Pin Count – “It’s the PIN’s Stupid – the PINS”.
- Improve Pin (ball) Utilization – Use better technology
- Vastly Increase Scalability, Availability, Performance

# What In Memory is Changing

- More On-Die processor Cache
  - Level Three Cache use and size is growing
- Total Memory Size Requirements Growing
- “Memory is Memory is Memory”
- Remote Memory – Make RDMA “Natural”
- Reliability of Memory
- Memory Technology is changing Getting smaller - nano
- Memory Speed is not tracking processor speed
  - Physical delays
- Layout constraints dominating designs
  - Signal Integrity
  - EMI
- Need Plug-n-Play
- Memory operations are for Cache Lines and Blocks not Bytes or bits

# Memory Hierarchy

- Level 0 – Processor Registers – access – xxx ps
- Level 1 – First level cache – I and D Cache – xxx ps
- Level 2 – Second level cache – unified 1- 8 MB – x ns
- Level 3 – Third level cache – General – IBM etc – x ns
- Level 4 – “Main Memory” – up to 1 TB – xx-xxx ns
- Level 5 – Virtual Memory – Paged Memory XTb – x us
- Level 6 – Fast Disk – 100GB – access x ms
- Level 7 – Active Storage – 2x growth from 8TB – xx ms
  - RAID and “Synchronous” backup - Failover
- Level 8 – Near Line Storage – from 96TB – xxx ms
  - RAID and Geographical distribution – Business continuance
- Level 9 – “ColdStor” – Petabyte increments – 30s
  - RAID, dump targets, large data bases, redundant
- Level 10 – “Tape Storage” Archive – Min, Hours, Days
  - Dumps, transport – questionable recovery

# Ideal Memory Interface

- Generic Memory Interconnect
- Extensible Design – Room for Options and Future Developments
- Byte Lane design - Support 1 to 64 Byte Lanes
- Does not change at each new Processor
- Plug'n'Play
- Self Timed From Interface – No distributed clocks
- Longer Interconnects - Extendable
- Possible switch or router Enabled
- Capable of RAIMM Operation (RAID without rotation)
  - Redundant Array of Inexpensive Memory Modules
- Transmission Path CRC32 Protected
- Smart Memory Module
  - Memory Caching Write and Read
  - Module Maintenance performed by Module
- Configuration and control via in-band messages

# Needed Connection Capabilities

- Addressing ***LARGE*** memory
- Variable Bandwidth Thin to Wide  
1-n “wires”
- Follow technology – 3Gb-> 6Gb->  
>10Gb-> + per second per wire  
pair
- Independent of actual Memory  
Technology
- Be independent of distance
- Scale 64 + bits of address
- Source and Destination  
identification
- Secure – Authentication and  
Encrypt
- Support redundancy
- Support MEMORY Coherency
- RDMA? Capability
- Scalable design
- Multiple implementation levels
- Higher Pin utilization
- Cost Effective
- Vendors make money
- Vendor differentiation
- Message optimized – not bit  
optimized
- Extended Memory + IP Capable
- Physical Layer Independent
- Multiple Physical Instantiations
- Common Protocol Layer

# Memory Vendor Design Freedom

- Define the interface tightly for commonality
- Plug-n-Play for user
- Implementation details to vendor
  - Tune for various price point – interface width
  - Tune for various reliability needs
  - Make available various sizes  $2^{10}$  –  $2^{64}$
  - Make available different power level
  - Make available different speed requirements
  - Change Technologies at will
  - Add Caching for better performance in interface
  - Add redundancy features – RAIMM techniques
  - Supply multiple technologies in same interface
  - Optimize yield
  - Differentiate and compete other than by cost
  - Meet Additional needs of systems vendors

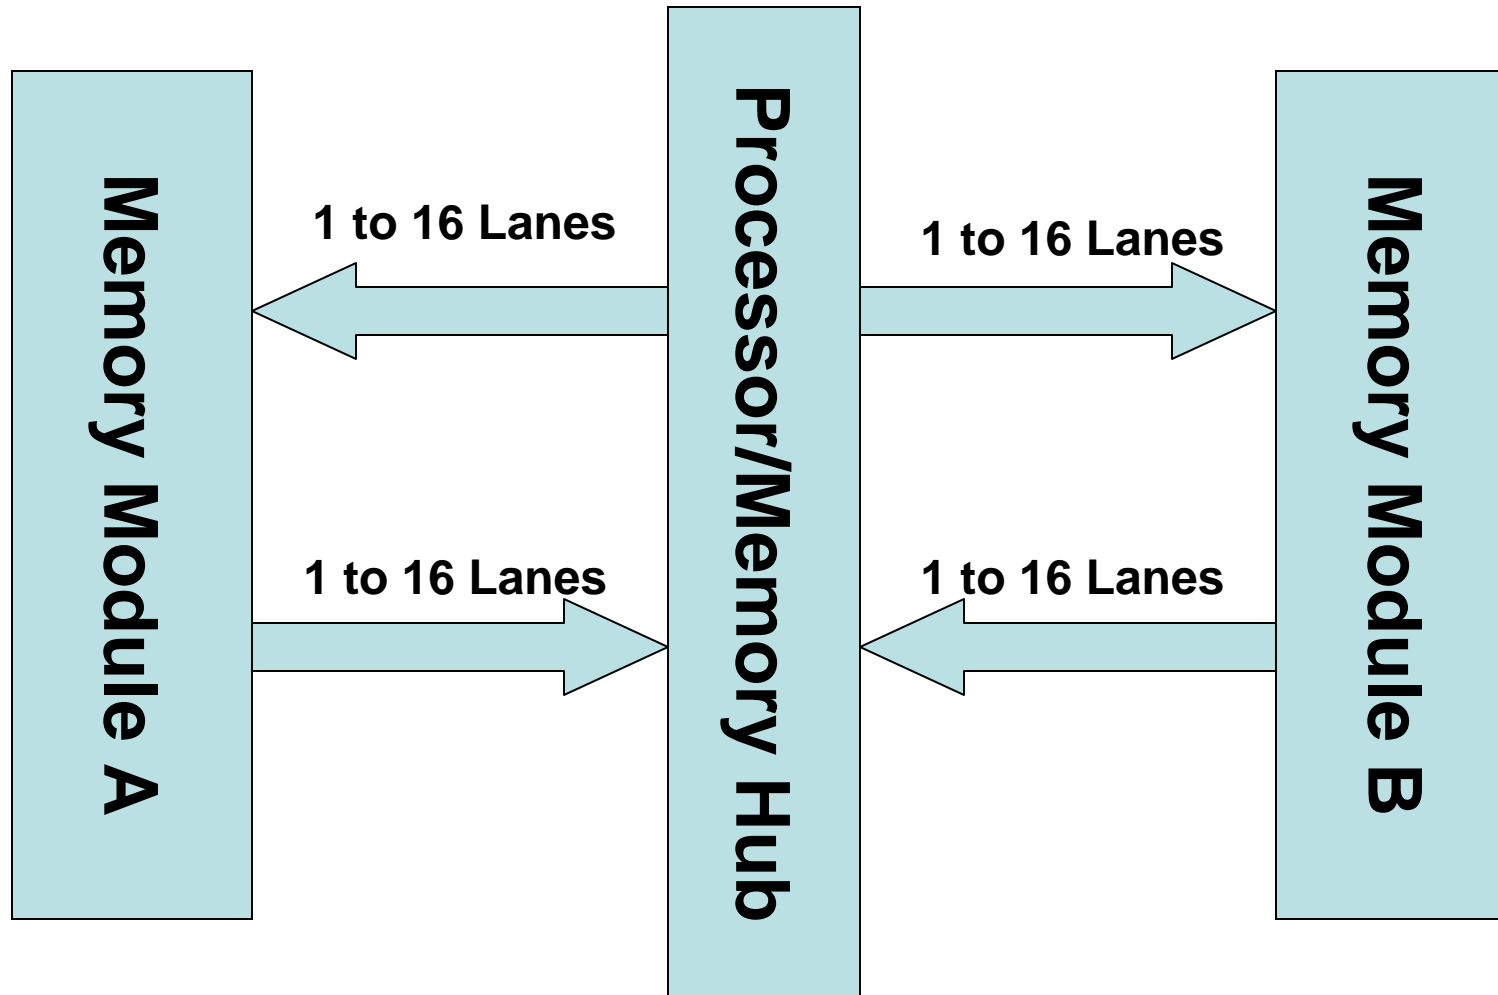
# Byte Lane Design Advantages

- Scalable in Width and Depth
- Capable of Redundant Memory Techniques
  - Redundant Array of Inexpensive Memory Modules
  - Support really paranoid customers
- Hot Replacement Possible
- Discovery and Configuration – In band
- Self Timed - Clock embedded in data
- Support orderly command structure
- Distance Limitations removed – except latency
- Reduce cost
- Well known SerDes designs
  - Ongoing development effort
- Possible switcher and router configurations

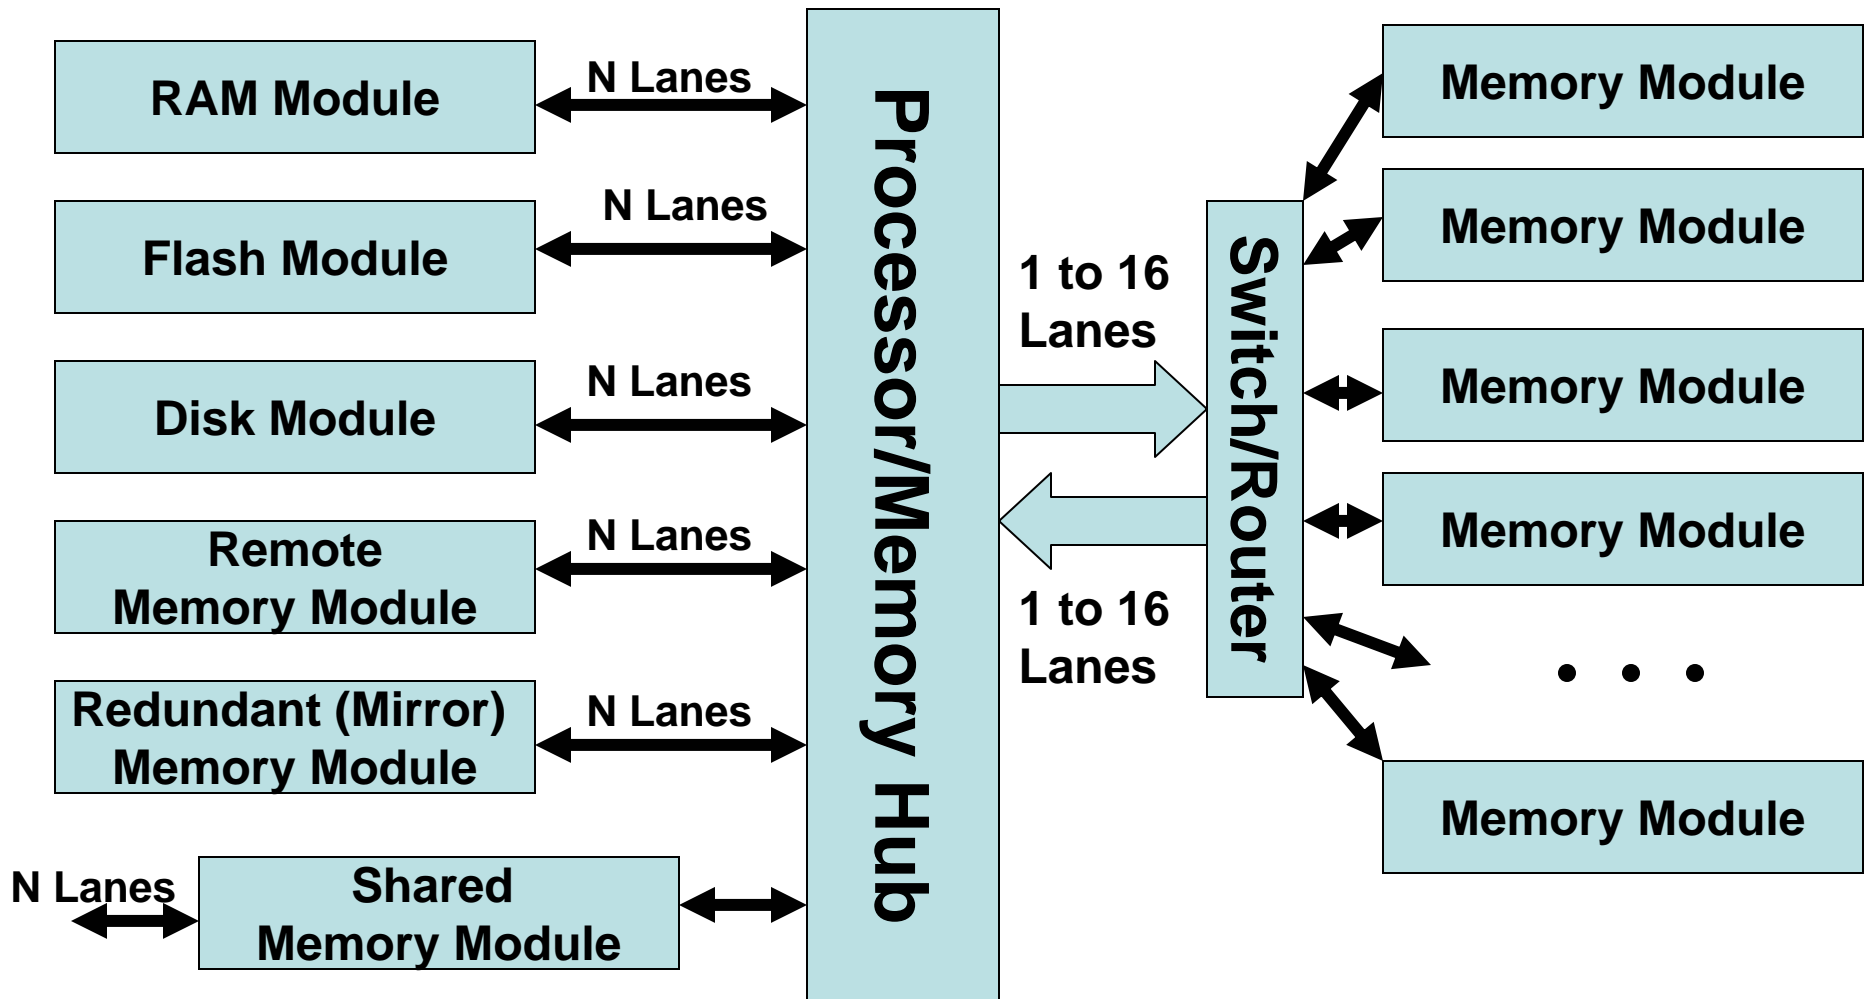
# Transfer Peak Bandwidth

	1 Lane 4 Wires	4 Lanes 16 Wires	16 Lanes 64 Wires	32 Lanes 128 Wires	64 Lanes 256 Wires
2.5 GBaud 2.0 Gb/s	<b>250 MB/S</b>	<b>1 GB/s</b>	<b>4.0 GB/s</b>	<b>8.0 GB/s</b>	<b>16.0 GB/s</b>
3.125 GBaud 2.5 Gb/s	<b>312.5 MB/s</b>	<b>1.25 GB/s</b>	<b>5.0 GB/s</b>	<b>10 GB/s</b>	<b>20 GB/s</b>
6.25 GBaud 5.0 Gb/s	<b>625 MB/s</b>	<b>2.5 GB/s</b>	<b>10 GB/s</b>	<b>20 GB/s</b>	<b>40 GB/s</b>
9.375 GBaud 7.5 Gb/s	<b>937.5 MB/s</b>	<b>3.75 GB/s</b>	<b>15 GB/s</b>	<b>30 GB/s</b>	<b>60 GB/s</b>
12.5 Gbaud 10 Gb/s	<b>1.25 GB/s</b>	<b>5.0 GB/s</b>	<b>20 GB/s</b>	<b>40 GB/s</b>	<b>80 GB/s</b>

# Simple Structure



# More Complex Structure



# Separate Protocol and Implementation Technology

- Establish Logical Protocol
  - Critical to Long-term success
  - Processor and Memory Independent
  - Packet Based Protocol most likely
  - Design to last 10+ years at compound growth
- Signaling Technology WILL change
  - Make use of developing techniques
  - Use cost driven technologies
  - Use performance driven technologies
  - Allow for future technologies

# Possible Near-term Solution Technology

- Extend Current RAS/CAS
  - Poor Pin utilization
  - Poor flexibility
- PCI Ex AS protocol
- Enet/IP based Protocol - RDMA
- InfiniBand
- SCI – Scalable Coherent Interface
- New Message Based Interface
  - Byte Lane optimized

# Packet Protocol

- Expandable - Extensible
- Multi-Generation Capable
- Not Specific to one type of Memory
- Write-only Packets
- Reliable
- Routable
- Best ROI
- Utilize Current Packet Technologies
- Allows unconstrained speed/distance tradeoff!
- Allows and Supports mixed solutions

# PCI Express AS Base

- Advantages
  - Routing switches being built
  - Is somewhat scalable
  - Low Cost implementations available
  - EUI identification built in.
- Disadvantages
  - Limitations of AS in large systems
  - Addressing poor beyond 64bits
  - Lacks configuration in peer to peer modes
  - Source routing –

# Beneficiaries of New Interface

- Processor Vendors
  - No needed knowledge of Memory Module Architecture
- Memory Vendors
  - More Optimized Memory Design – Less physical driving constraints
  - Add interface to chip for smallest systems
- System Builders
  - More Options for Memory
  - Better Signal Integrity design
  - Faster Time To Market
  - Capable of UPGrade after sale
  - Expanded System Architecture to meet specific design goals
  - Component Compatibility over time for longer product life
- Memory Module Vendors
  - More design freedom in sub-unit design
  - Opportunity to add capabilities for differentiation
  - Memory is Memory is Memory –different technologies
- Users
  - Plug-n-Play
  - Capable of upgrading system memory and performance
  - Ability to support special applications needs.

# Why MSC?

- Microprocessor Standards Committee
  - Sponsors Bus Structures
  - Sponsors Microprocessor Related Projects
  - Has Expertise
  - Developed SCI and derivatives
  - Developed IEEE 1394 (Firewire, iLink, etc)
  - Microprocessors need this development

# Next Steps

- Gather the Experts – a Study Group
- Define Goals – Request PAR
- Define Long Range Architecture
- Define A Fully Scalable Protocol
- Learn From/Incorporate Principals From
  - SCI, IB, IP, etc
- Make use of other Developed Standards
- Plan for expansion with Options
- Put in Place by 2005/6/7

# Conclusions – Final Plea

- Break the tradition of the IBM PC
- Don't bit optimize as first principle
- Make a Long Lived Generic Memory Interface
- Make an Interface that can be implemented over a wide variety of applications
- Expand the scope to larger Memory objects
- Shift to a message style packet based Interface
- Move care of the Memory to the Memory
- Use Optional Fields to build in very dumb to very capable module capabilities
- Remove Distance limitations
- Allow authentication of module and Redundancy
- Remember "Memory is Memory is Memory"