SiliconGraphics

Building High Performance Server Network Interfaces

Joe Gervais

Product Line Manager

High Performance Networking

Interconnect Hierarchy

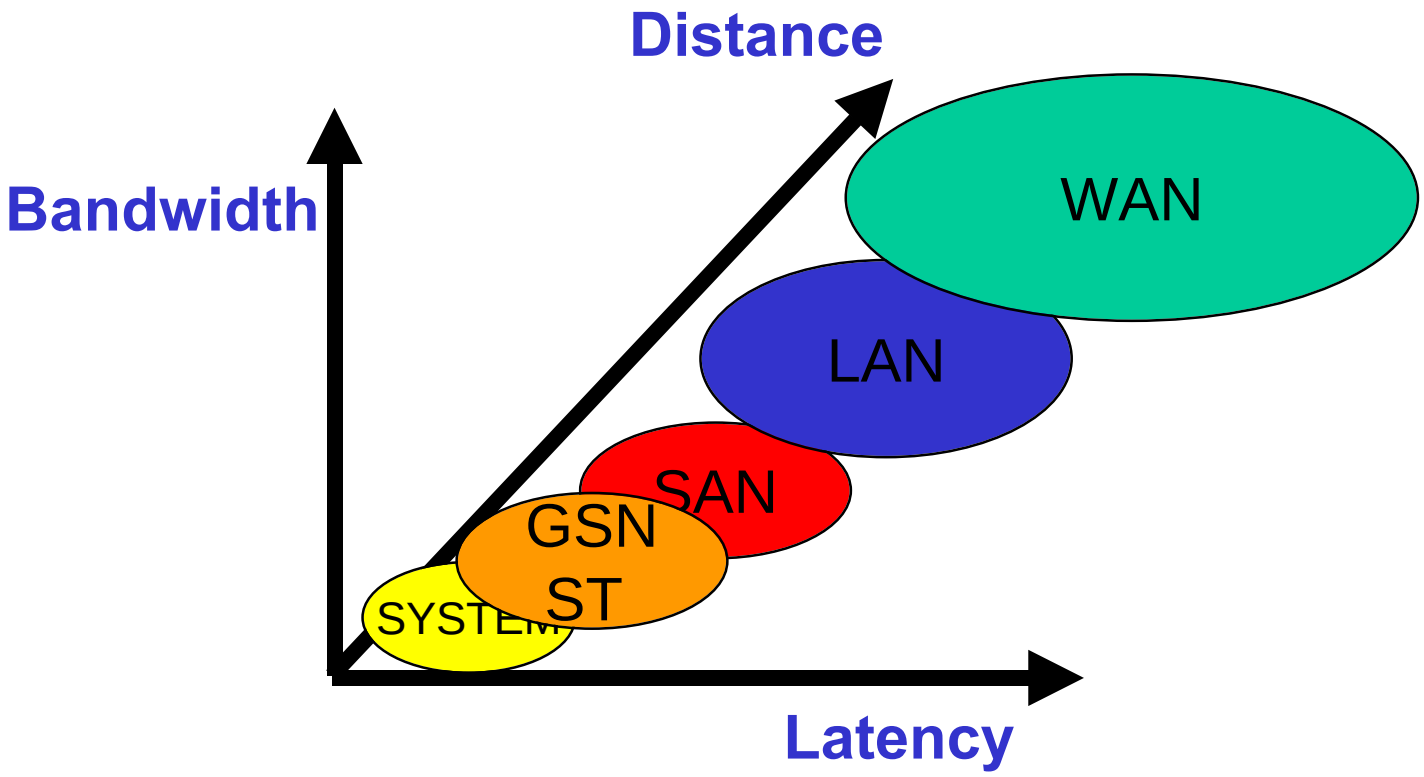

Gigabyte System Network™, or GSN™

- Is the highest bandwidth and lowest latency interconnect standard
- Provides full duplex dual, 6400 megabit (800 megabyte) per second of error-free, flow controlled data
- Provides for interoperability with Ethernet, Fibre Channel, ATM, HIPPI-800 and other standards

GSN Topology

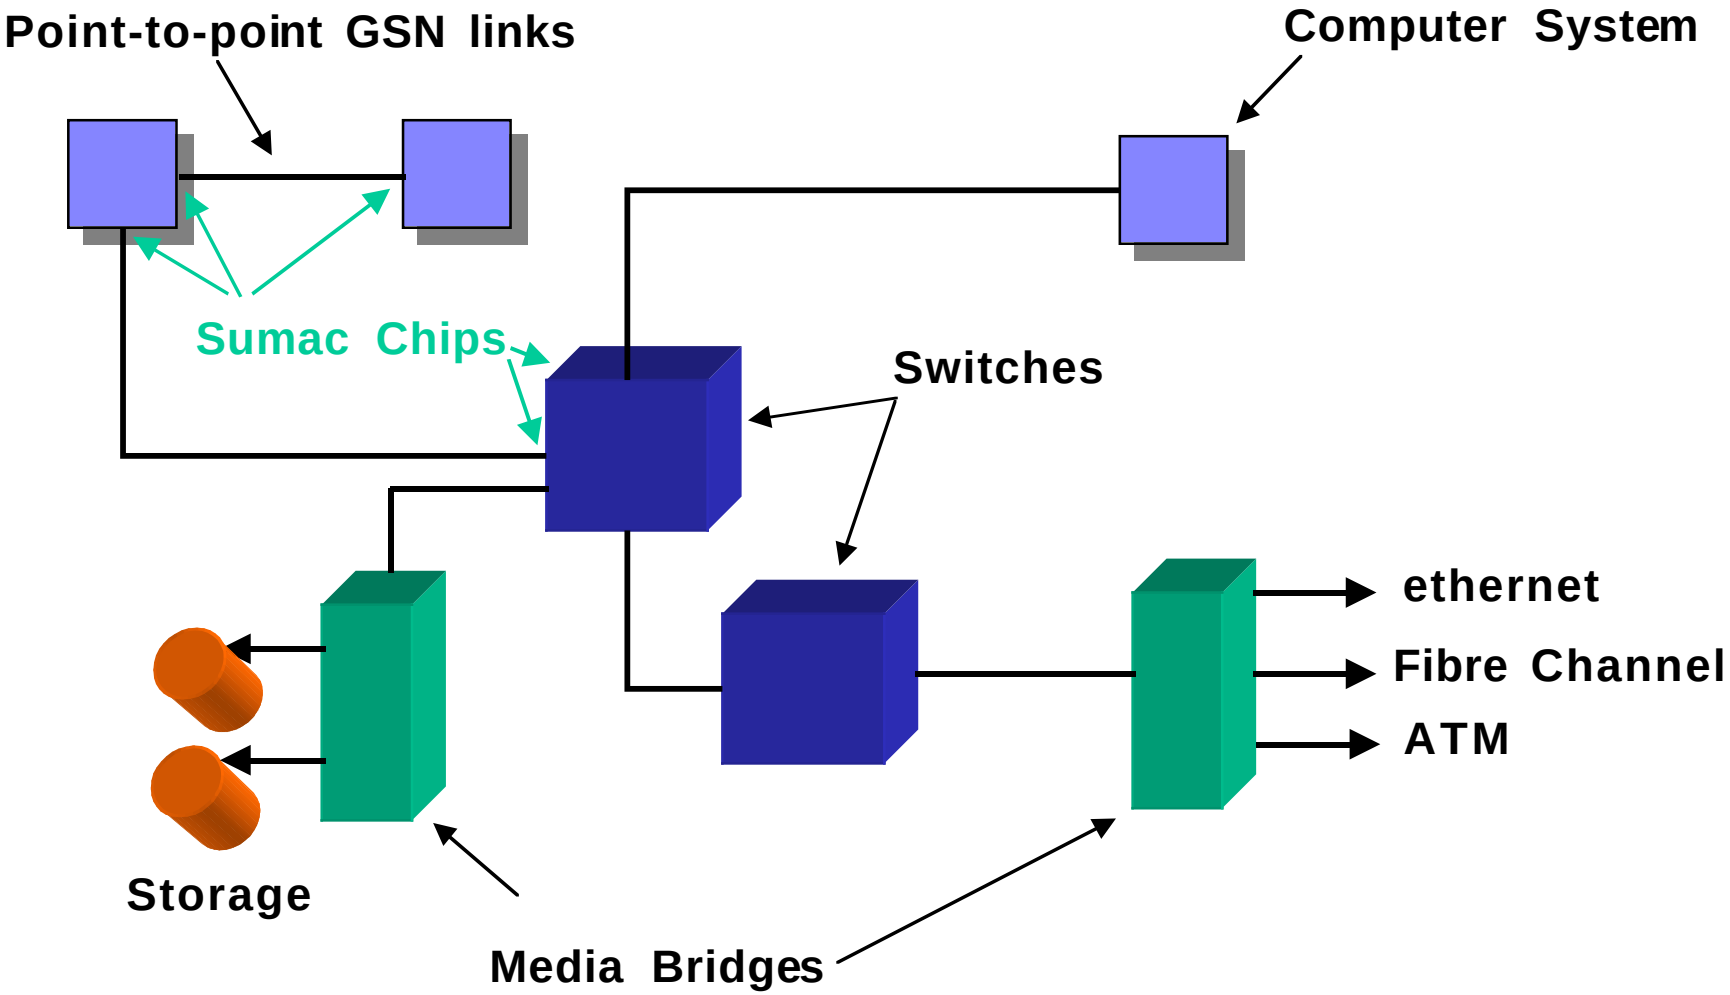

GSN Link Features

SiliconGraphics

SuMAC vlsi chip

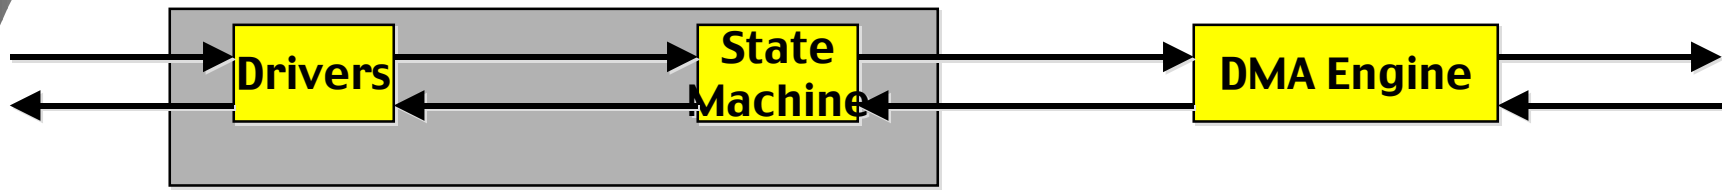

Physical

10 Gbit/s
- 6.4 data
- 1.6 control
- 2.0 encoding

Copper
- 40 meters
- 20 bits wide @500MHz

Fiber
- 1 Km
- 10 Bits wide @1GHz

Logical

Multiplexed
- 32 Byte cells
- 4 VCs

Error Control
- Link CRC
- End to End CRC
- Go back N retry

Flow Control
- credits for each VC

Data Buffers
- on chip

Messaging

IEEE 802 MAC Header
- expanded length field

Protocol

Any IEEE defined ULP
- TCP/IP
- Schedule Transfer

GSN
Gigabyte System Network

SiliconGraphics

The technology is ideal wherever organizations require timely movement of large amounts of information including

- **scientific and technical computing**
- **HDTV**
- **data mining**
- **transaction processing**
- **video and film archiving**
- **storage management**

SGI XIO GSN Adapter

Full bandwidth, high performance adapter

Performance (Beta Software)

- Raw GSN 796 MB/s, 791 MB/s bidirectional
- ST 560 MB/s, 769 MB/s memory striping
  - **Uses 5-8% of one CPU**
- TCP 280 MB/s, UDP 500 MB/s
- BDS/ST 450 MB/s, 690 MB/s mem. striping
- One-way latency 4-9 usec (mem-to-mem)
  - 4 - 6 usec pt - pt 5m cable
  - 6 - 9 usec **50m cable plus switch**
- Packet transmit rate 1.45 million PPS
- ST scaling: aggregate bandwidth scales at 98% of linear from 1-3 NICs

… Stay Tuned

SGI XIO GSN Adapter

- *Powered on August 1998*
- *Interoperated with Essential Switch September 1998*
- *To be offered in single and dual XIO versions*

SGI XIO GSN Adapter

To be offered on Origin 200 GigaChannel, Origin 2000 and Onyx2

- Design has very efficient CPU utilization
- Memory a different story

SGI XIO GSN Adapter

Origin 200 GigaChannel, Onyx2 Deskside

- Expected to not exceed 400 MB/s total

Origin 2000 Deskside, Onyx2 Rack (8P)

- Expected to not exceed 500 MB/s, FDX

Origin 2000

- 12 CPU and larger should be able to achieve full bandwidth for proper apps

SGI XIO GSN Adapter

Bandwidth requirements are very memory intensive

- **Node Boards tested to 504 MB/s memory bandwidth to GSN boards**
- **6 node boards for full bandwidth**
- **Most applications will only be able to take advantage of ~ 500 MB/s performance**

Adapter Architecture

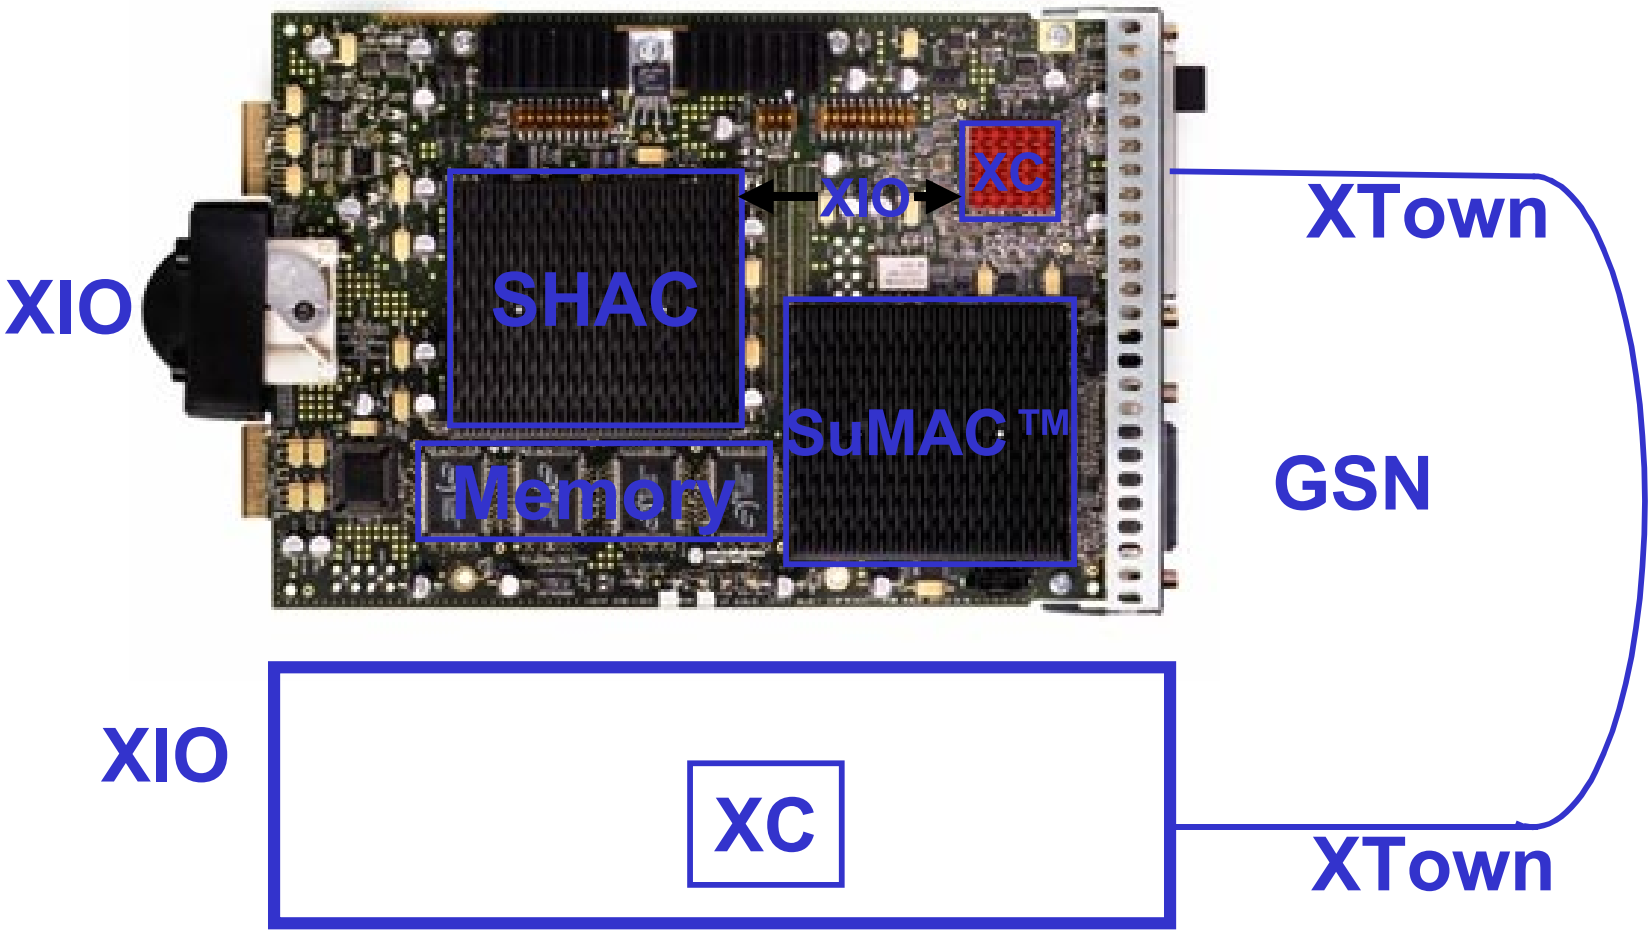

Bus Choices

- 64 Bit / 33 MHz PCI - ~250 MB/s Total
- 64 Bit / 66 MHz PCI - ~500 MB/s Total
- SGI Dual XIO GSN Adapter hardware tested to 99% of line rate
  - 791 MB/s per direction
  - Software overhead will reduce actual performance

Latency

- HIPPI hardware latency with XIO adapter on the order of 80-90 u sec
  - MPI adds ~ 30 u sec
- GSN hardware latency expected to be on the order of 6 u sec or lower through a switch
  - MPI may be as low as 30 u sec total for a 4x decrease

SGI SHAC ASIC

- DMA Engines per VC
- ST Data Transfer Engine
- TCP Checksum support
- Supports Multiple Interrupt Queues
  - true parallel driver architecture - not just multi-threaded