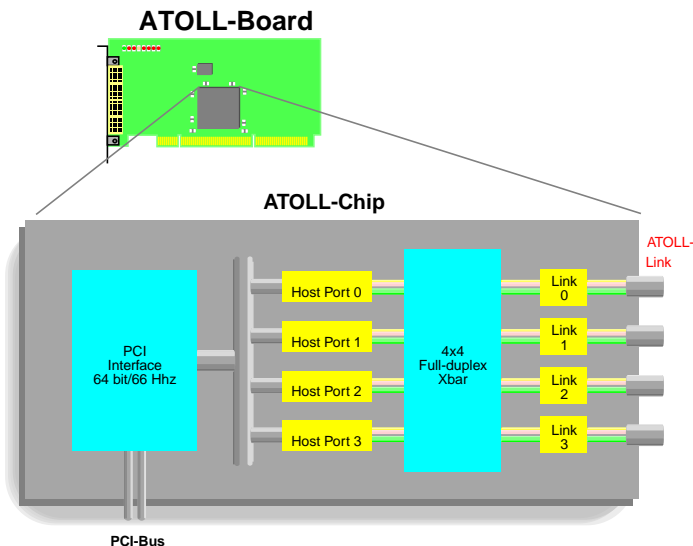
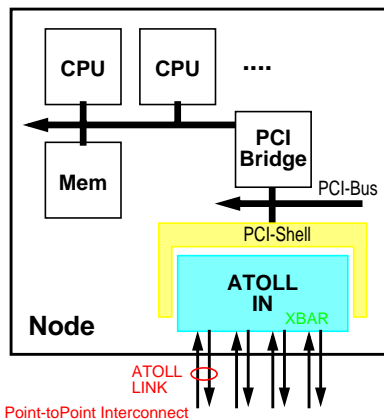


the ATOLL system

atomic low latency



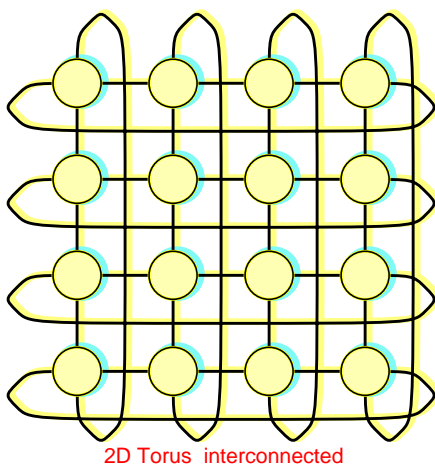
Processing Node with ATOLL-Board



$2 \times 4 \times 250\text{MBytes/s} = 2\text{GBytes/s}$

System with 16 Nodes

extendable to $n \times m$ Nodes



2D Torus interconnect

The **ATOLL-Cluster Interconnect Network Board** provides a simple and efficient way to build **System-Area-Networks (SAN)**. ATOLL is the enabling technology to form cost-effective and high-performance Cluster of Workstations or PCs based on commodity-off-the-shelf processing nodes. The low latency for communication start-up and the hardware support for message handling deliver new levels of performance for parallel processing based on message passing. The four interconnections of the ATOLL links offer the ability to construct arbitrary network topologies without additional switching hardware. The typical topology will be a 2D-grid or torus.

The highly integrated CMOS ATOLL-Chip is the basic building block for an ATOLL-SAN and it provides the following features:

- Advanced ASIC technology: about 4.5 million transistors in a state of the art IC technology (0.18um CMOS), running at 250MHz
- **four independent network devices** (host ports) linked through an internal **4x4 bidirectional crossbar** to four byte-wide links to the interconnection network
- four optimized host interfaces to the driver software for the simultaneous access of four processes or threads to the ATOLL network
- routing is performed by source path wormhole routing (27 clock tics of 4ns per routing stage = 108ns!)
- **high interconnection bandwidth** (~500MB/s per bidirectional link; **2GBytes/s per chip**)
- minimal hardware **message latency of ~1,2 μs**
- user-space to user-space communication without involvement of the operating system
- **reliable data transmission** based on hardware support for error detection and correction and reverse flow control of messages which eliminates the need of buffering on the sender side and provides a significant reduction of protocol overhead

The PCI-X-Standard Bus System (32/64bit, 66-133MHz) as host interface allows the utilization of all processing nodes supporting the PCI-X-Standard. The **PCI-X-Interface** can transmit and receive data at a rate of **1064 MBytes/s** and supports PIO- and DMA-Modes.

Message Passing is supported as the basic programming model. The message passing library **MPI** is provided using the ATOLL communication API PALMS.

The ATOLL-Chip, package, board and system software is a development of the University of Mannheim, Chair of Computer Architecture.