Multi-gigabit Switching and Routing

Gignet '97 — Europe: June 12, 1997.

Nick McKeown
Assistant Professor of Electrical Engineering and Computer Science

nickm@ee.stanford.edu
http://ee.stanford.edu/~nickm

Multigigabit and Terabits

1. How Routers have Evolved
2. Multigigabit Routing
3. Terabit Switching and Routing

The Evolution of Routers

The first shared memory routers

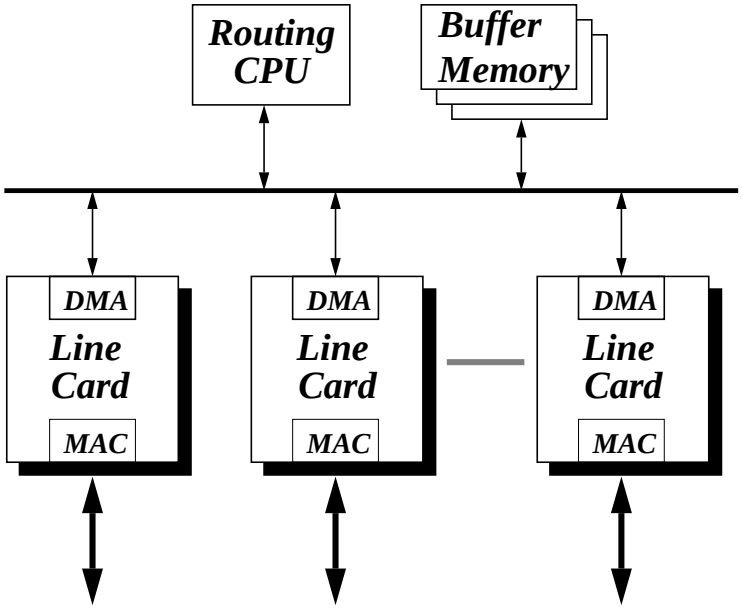

The Evolution of Routers

The first shared memory routers

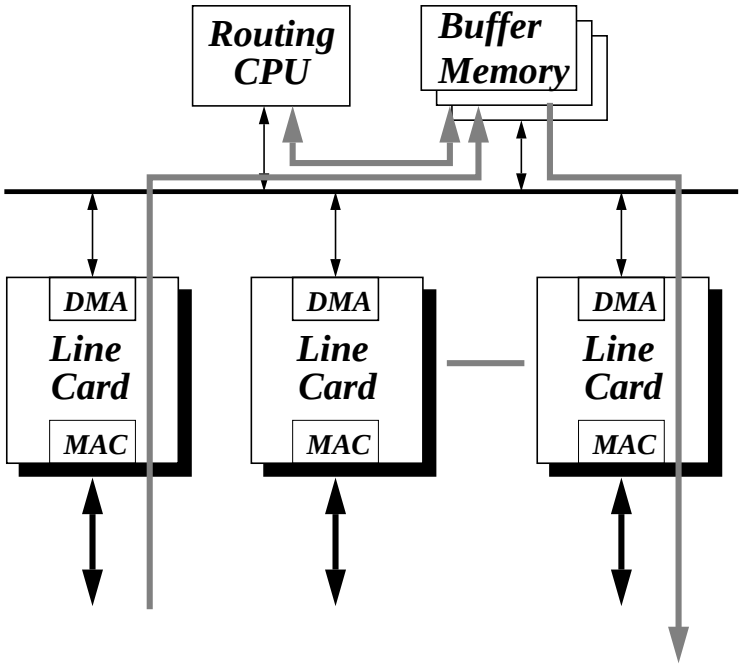

The Evolution of Routers

Reducing the number of bus copies

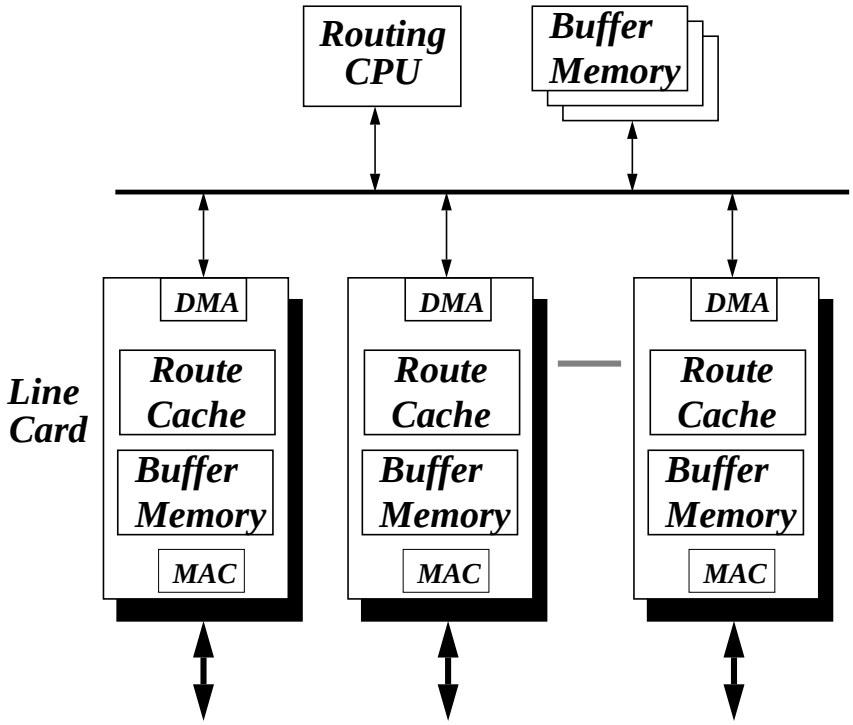

The Evolution of Routers

Reducing the number of bus copies

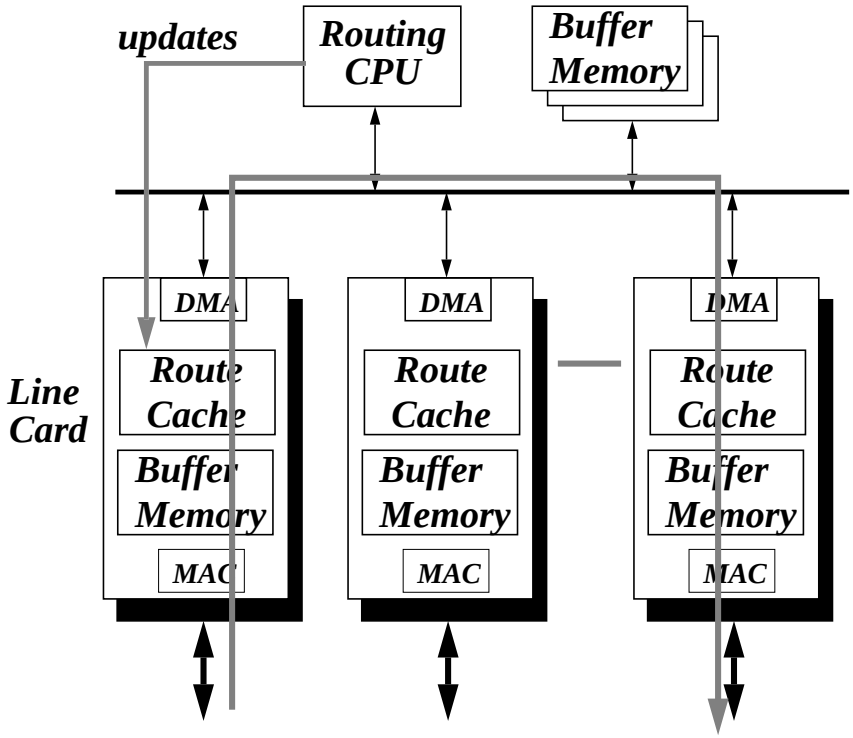

The Evolution of Routers

Avoiding bus contention

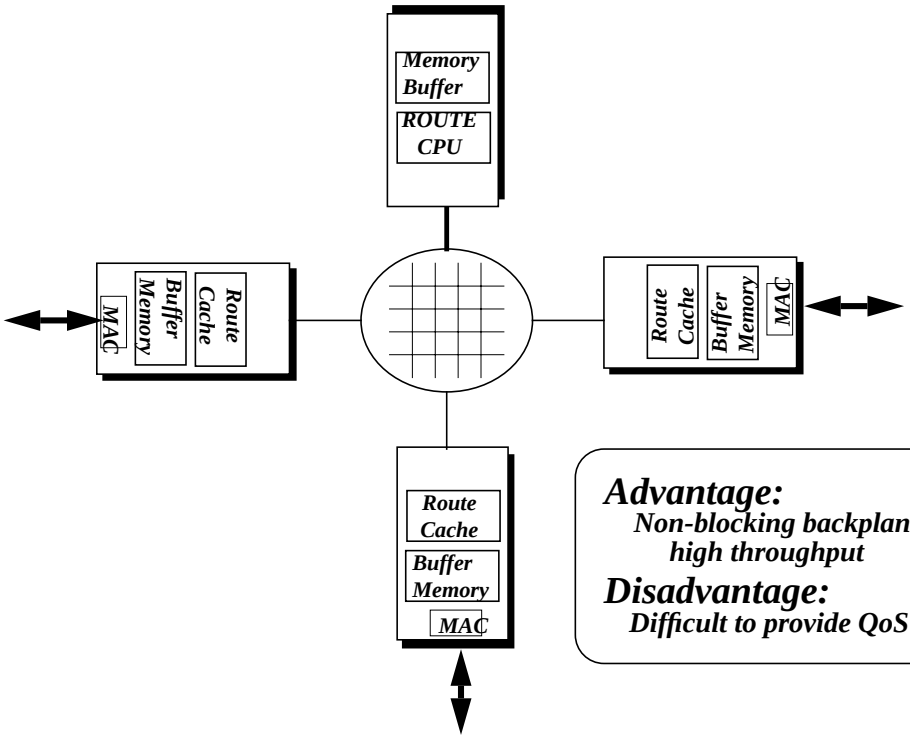

Advantage:
Non-blocking backplane—high throughput

Disadvantage:
Difficult to provide QoS

Multigigabit and Terabits

1. How Routers have Evolved
2. Multigigabit Routing
3. Terabit Switching and Routing

Multigigabit Routing

BBN's Multigigabit Router

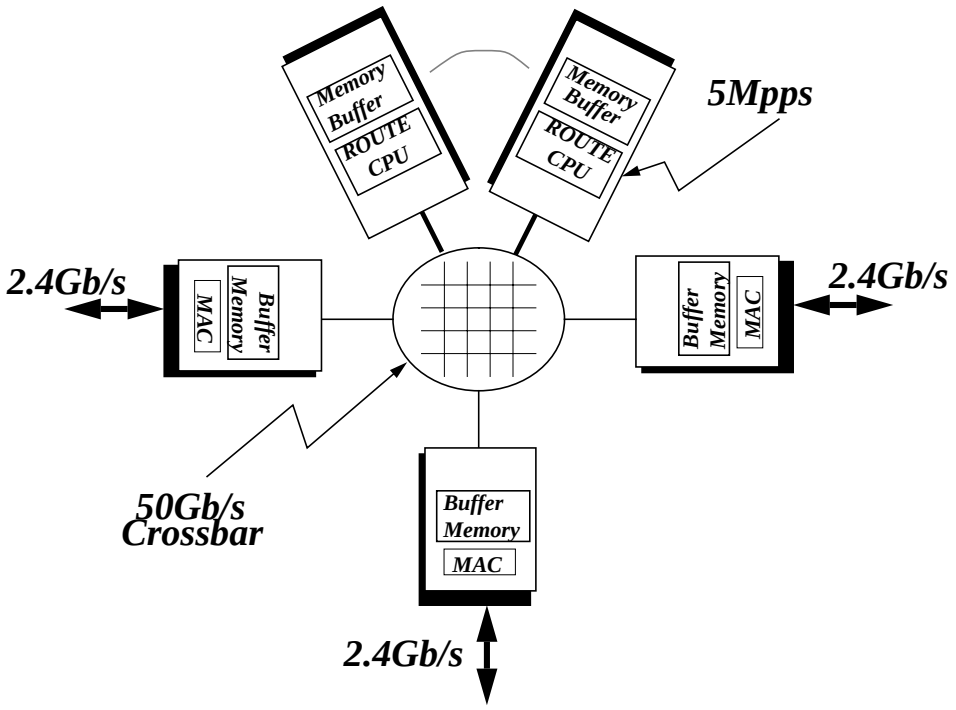

Routing Lookups

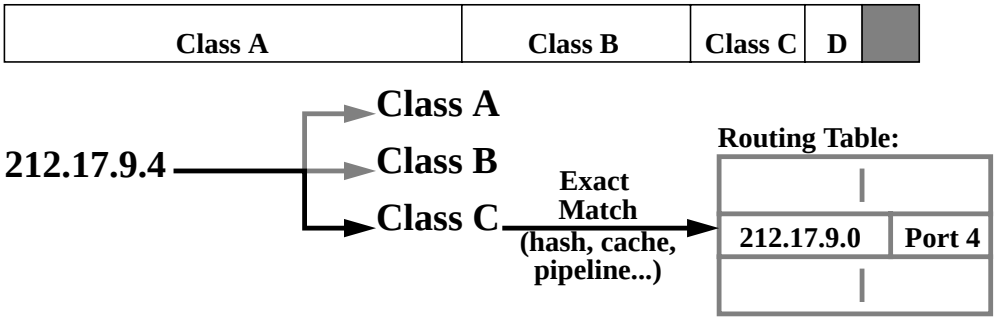

Routing Lookups with CIDR ("supernetting")

CIDR uses "longest matching prefix" routing:

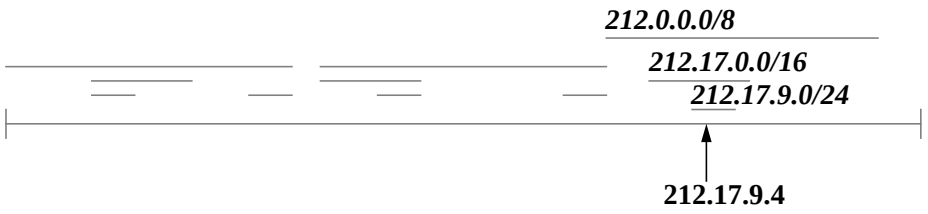

Hashing, caching and pipelining are hard!

Solution 1: *Label Swapping*

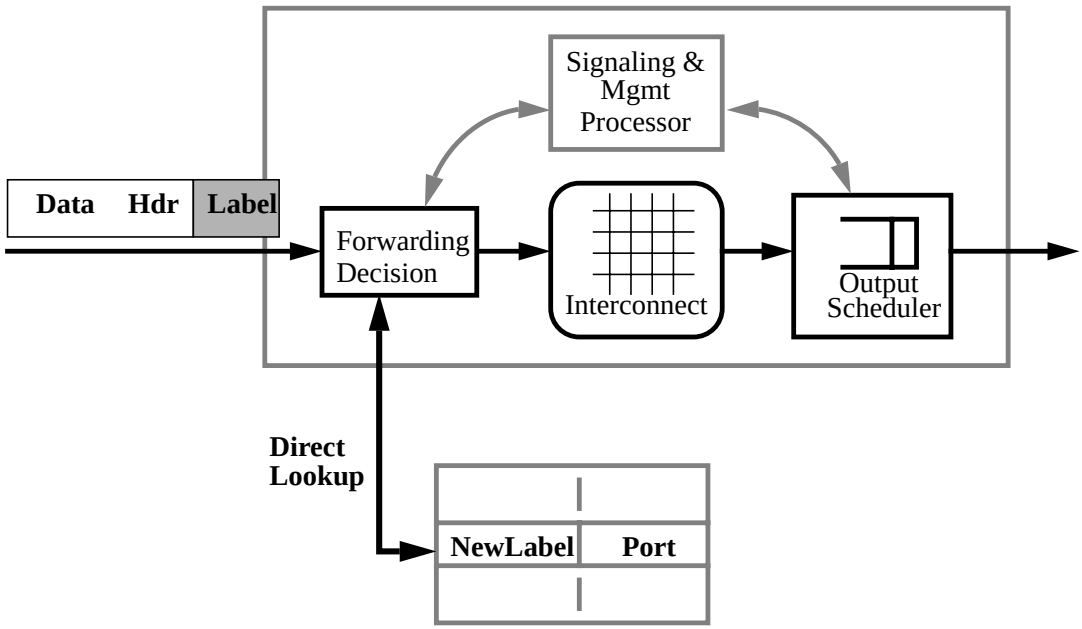

IP Switching, Tag Switching, ARIS, Cell-switched Router,....

Solution 2:
Perform Lookups Faster!

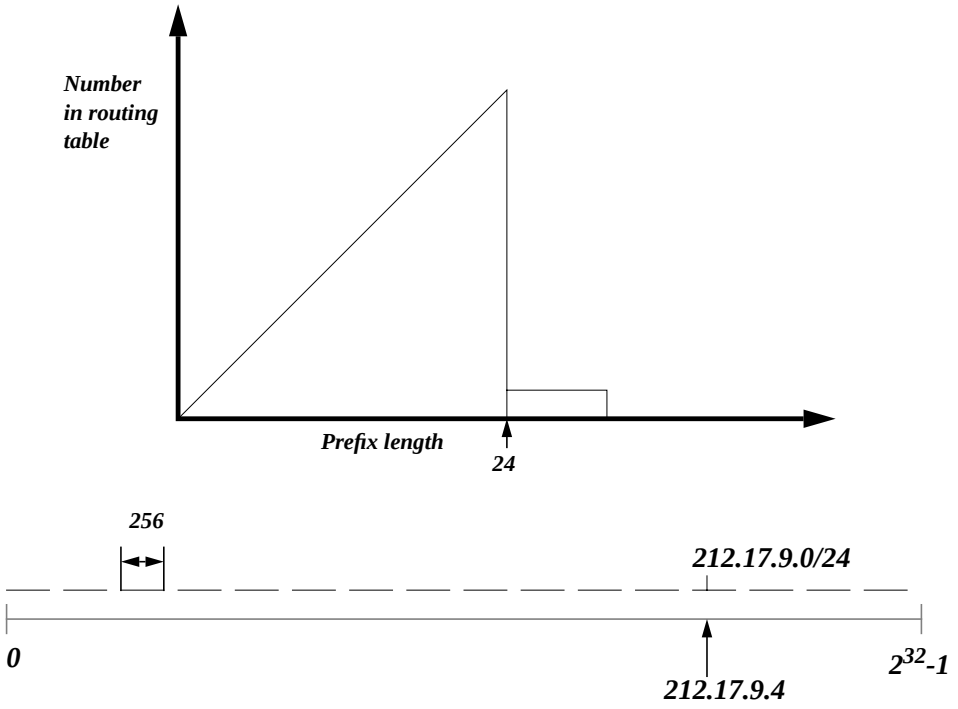

Solution 2 (cont):
20 million lookups per second

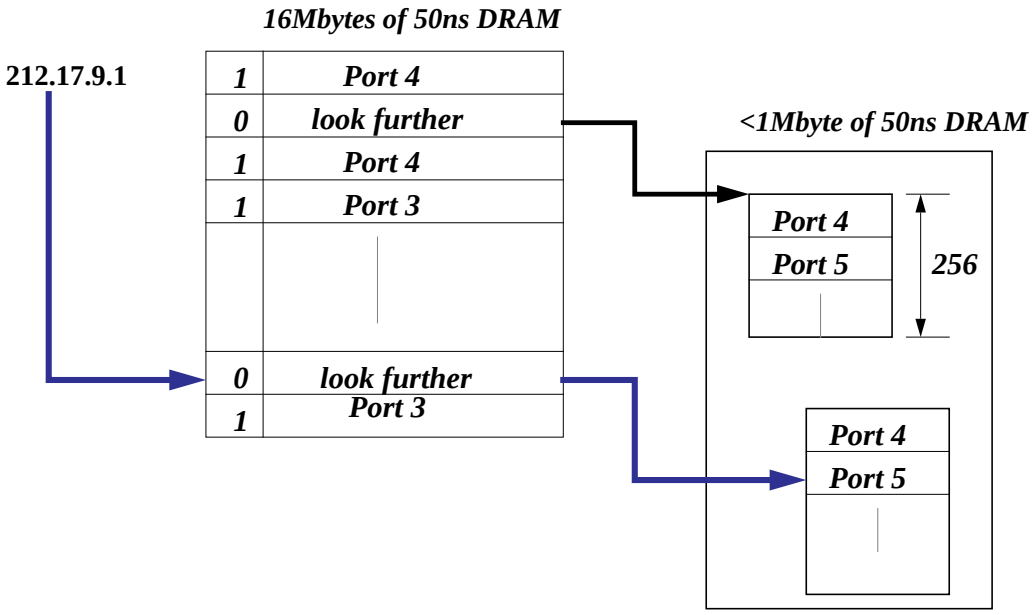

Multigigabit and Terabits

1. How Routers have Evolved
2. Multigigabit Routing
3. Terabit Switching and Routing

The Tiny Tera

"Soda can" switch core
32x32 switch, ~16Gbps per port
Aggregate bandwidth: 0.5Tbps

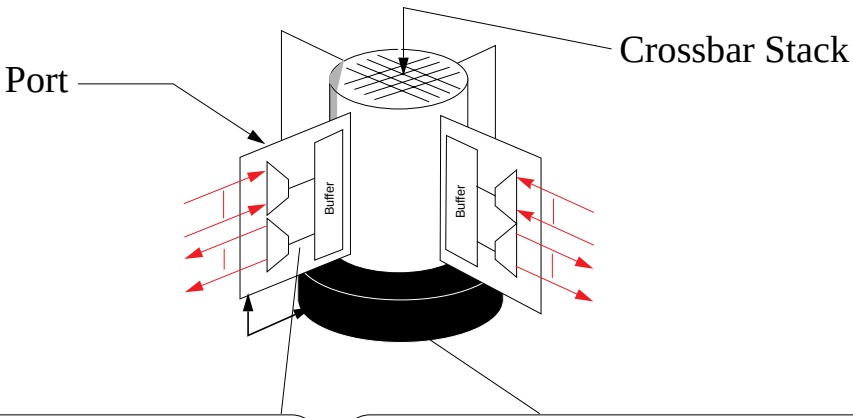

Requires high speed chip-to-chip links.

Schedulers must be fast, fair and efficient.

High Bandwidth Parallel Datapath

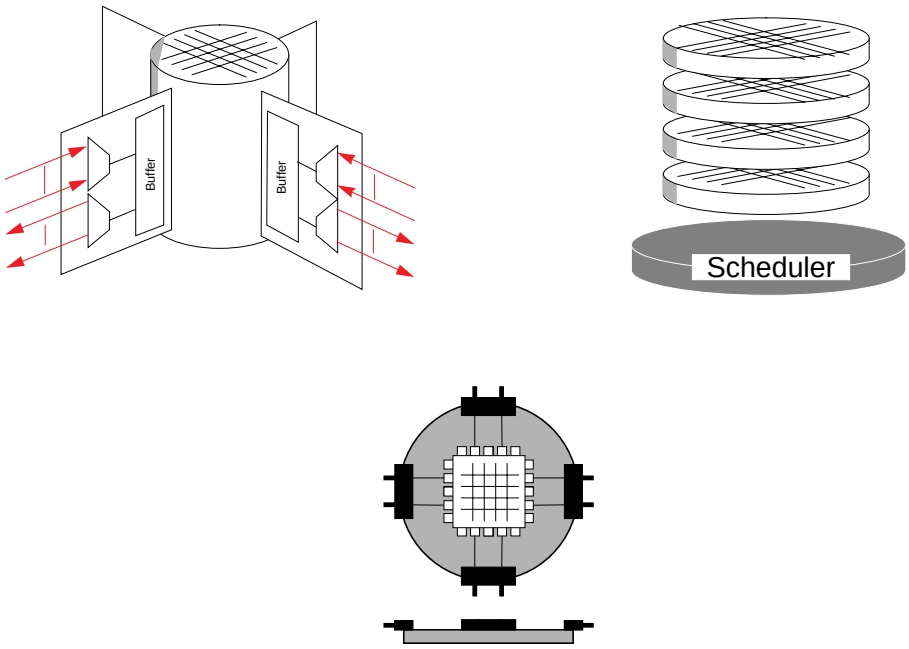

The Tiny Tera Port Architecture

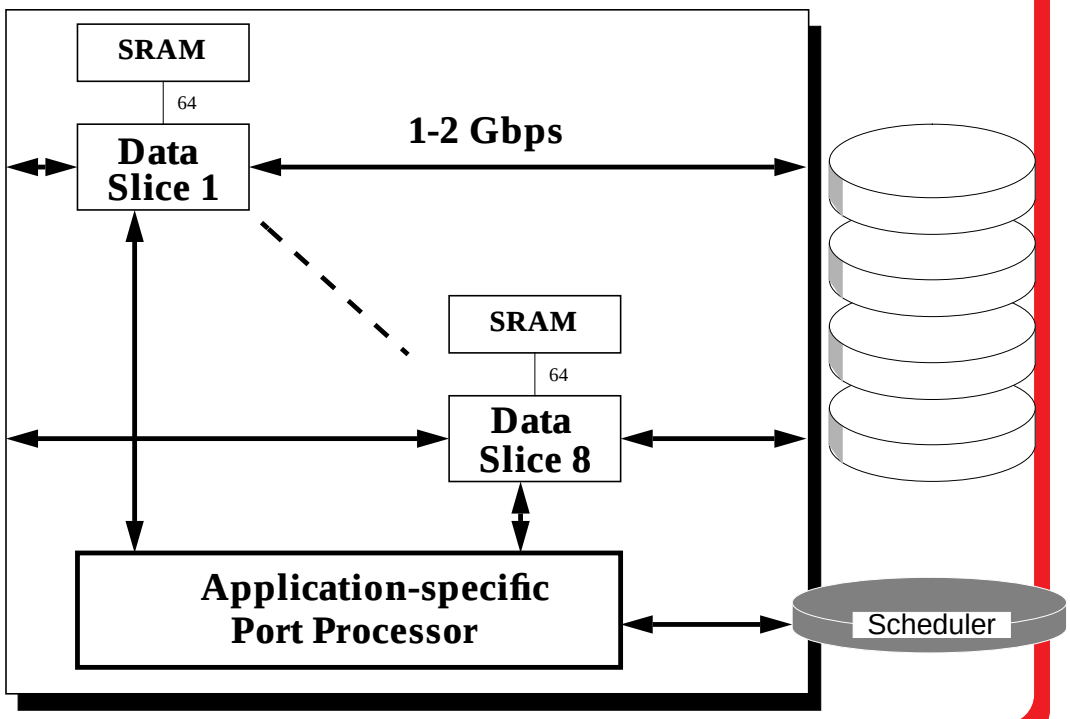

The Tiny Tera

http://tiny-tera.stanford.edu/tiny-tera/

32 ports, 16 Gb/s per port.
***Input*-queued architecture.**
High bandwidth *parallel* datapath.
Efficient unicast *and* multicast.
Four priority levels.
Fixed *and* variable length packets.
Tag switching.
***Asymmetric* high-speed serial links.**