Lara Networks

Technology Solutions for a Better Internet™

Lara Networks

Background

- Dr. David V. James
  dvj@alum.mit.edu
- Editor or Chair of
  - 1596, 1596.3, 1596.5, 1212, LVDS
  - Low latency fairness
  - Fast standard=>open to change
- Why networks
  - SAN, MAN, WAN, LAN leverage
  - Backplanes limited by “Moore’s Law”

Lara Networks

From a bus perspective...

- Ethernet without the bus constraints
  - Low latency, guaranteed bandwidth
  - Fairness with bounded latency
  - Accurate wallclock synchronization
- Well written, is out there...
  - IEEE 1596-1992 SCI
  - IEEE 1394-1994
  - Written&posted pseudo-draft
- Special interests
  - Deterministic scheduling in the home
  - Simple/sufficient QOS on the first mile

Lessons of the past...

- Flow control mandates 2-out-of-3
  - Low latency transmissions
  - Fair bandwidth allocation
  - High bandwidth utilization
- Feedback control systems
  - Low latency signaling
  - Control passes asynchronous packets
  - Separate synchronous queues
- Other observations
  - Local control => global perversions
  - Fairness is inherently “approximate”
  - Strange beating sequences DO OCCUR

Arbitration classes

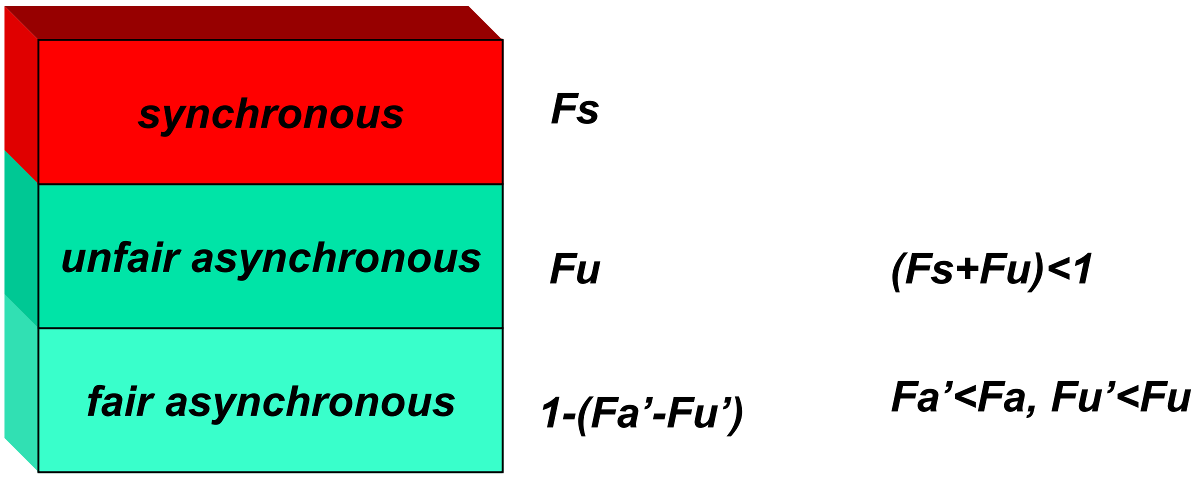

- Low latency & guaranteed BW
- Bounded latency & guaranteed BW
- Fairness on residual bandwidth (unused as well as nonprovisioned)

Lara Networks

Physical layer assumptions

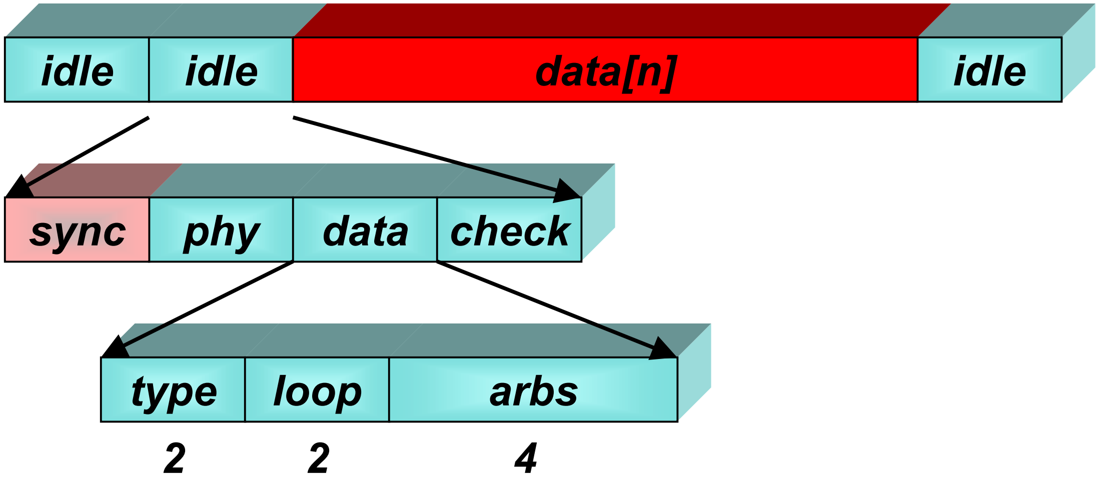

- Packets framed by idles
- Idles have control information
- Arbitration uses 4 arbs bits

Lara Networks

Arbitration related components

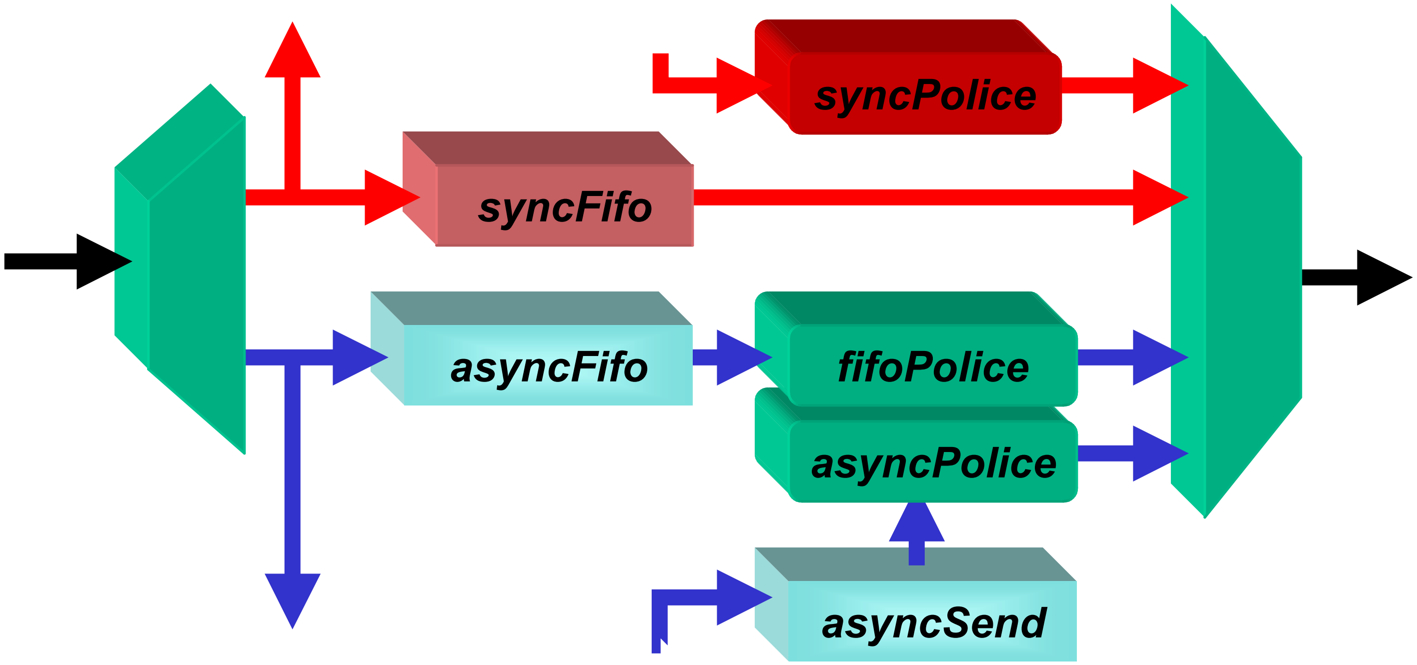

- Distinct sync and async paths
- Load dependent policing

Lara Networks

Synchronous-send policing

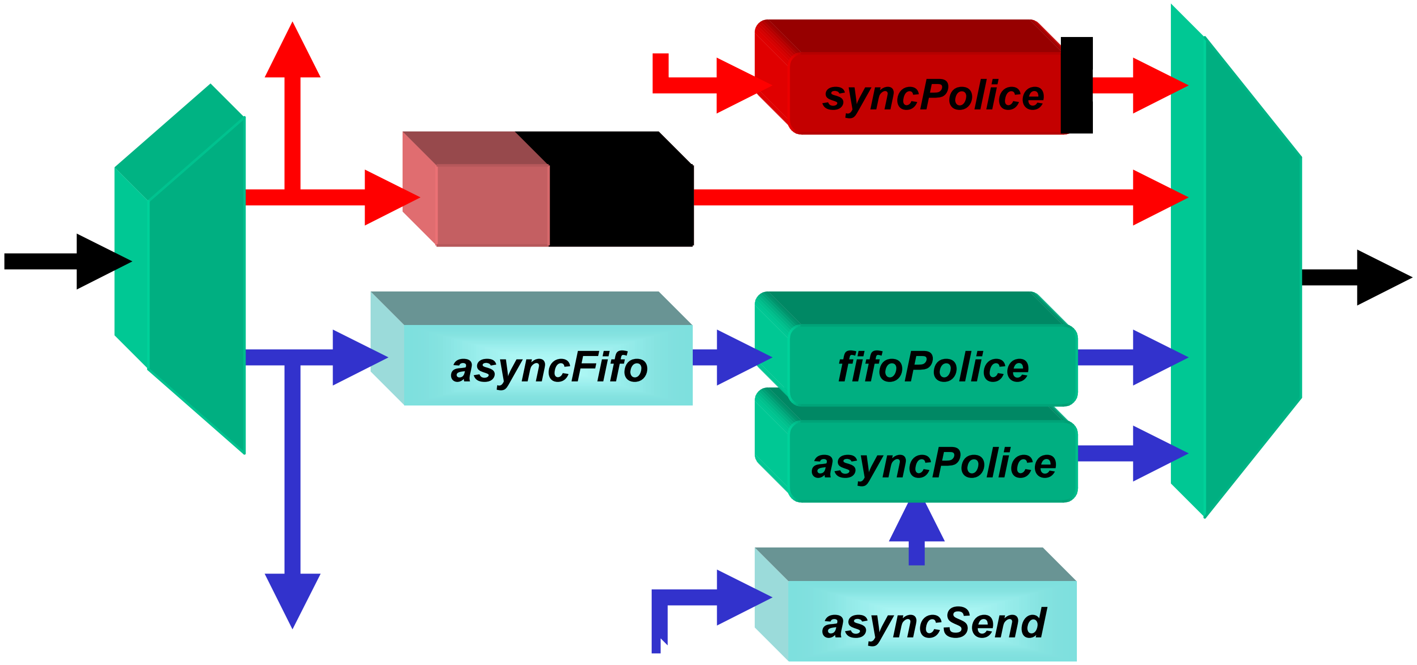

- FIFO overflow avoidance
- Rate limiting and spreading

Lara Networks

Synchronous backlog policing

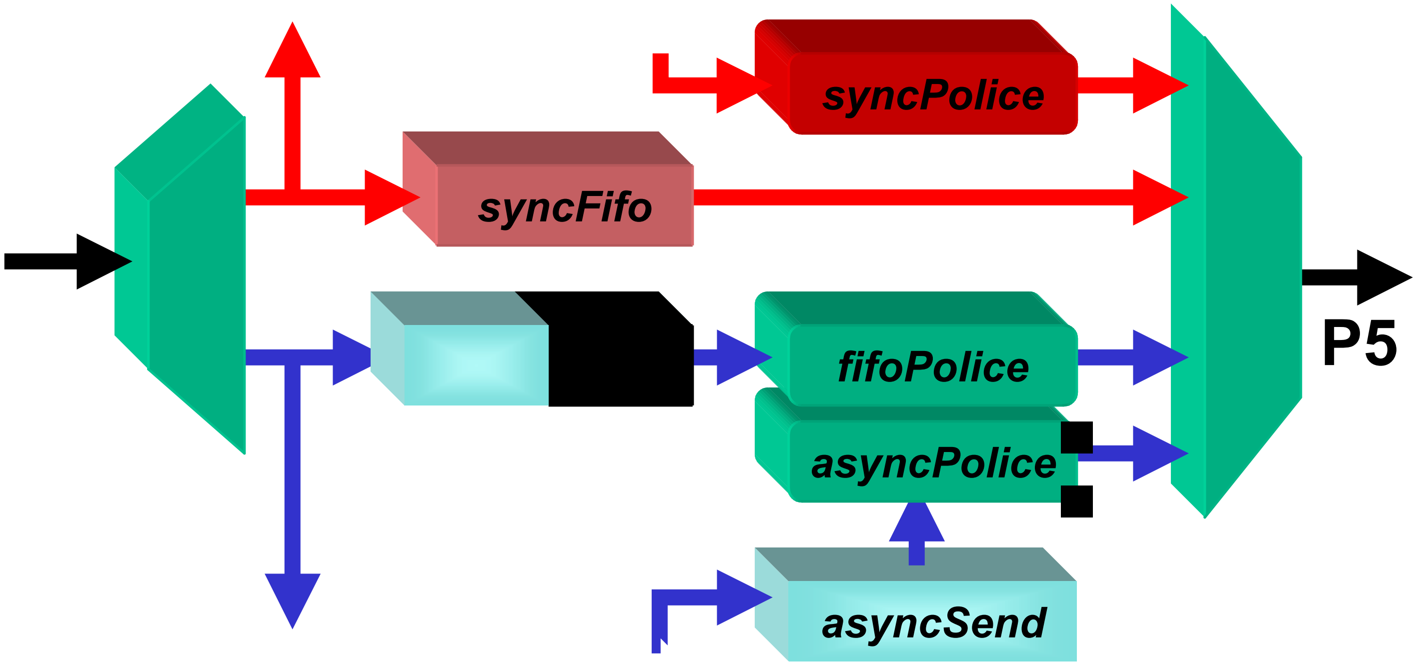

- Avoid future reserves consumption
- Limit own asynchronous, signal others

Lara Networks

Synchronous backlog assistance

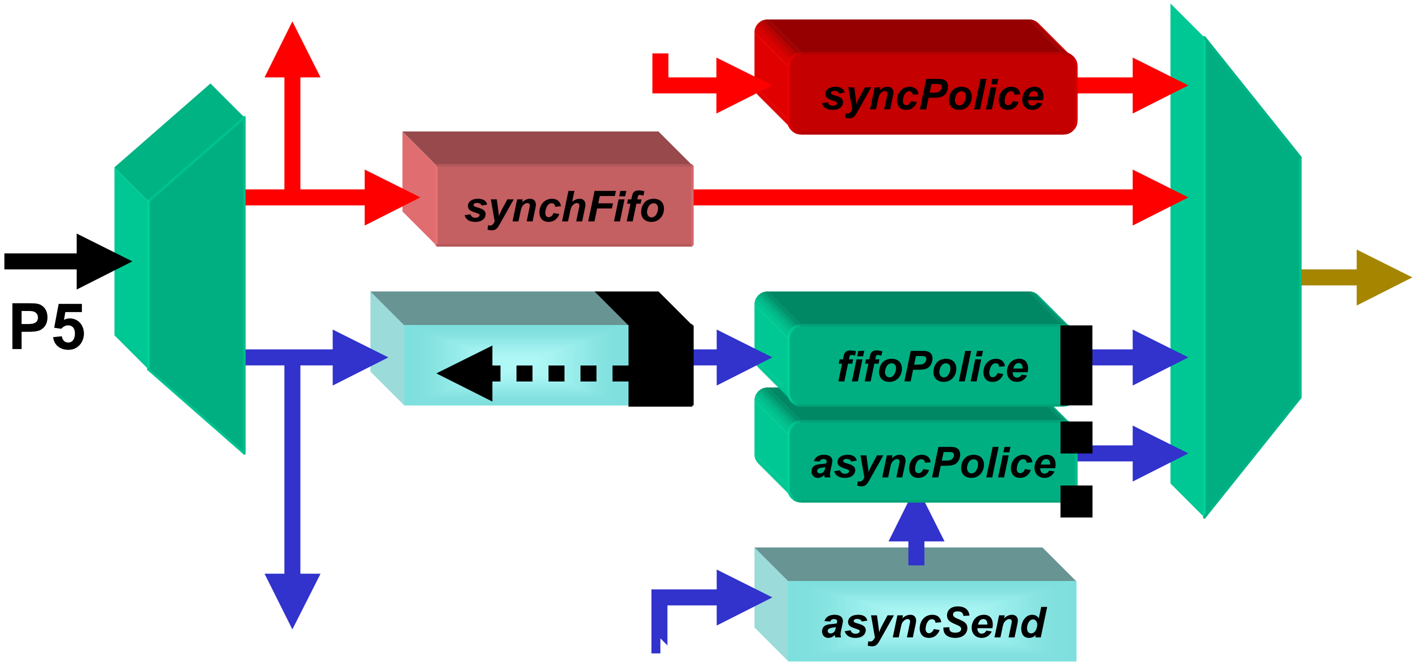

- Sync assistance => asyncFifo filling
- Throttle asynchronous insertions

Fairness policing

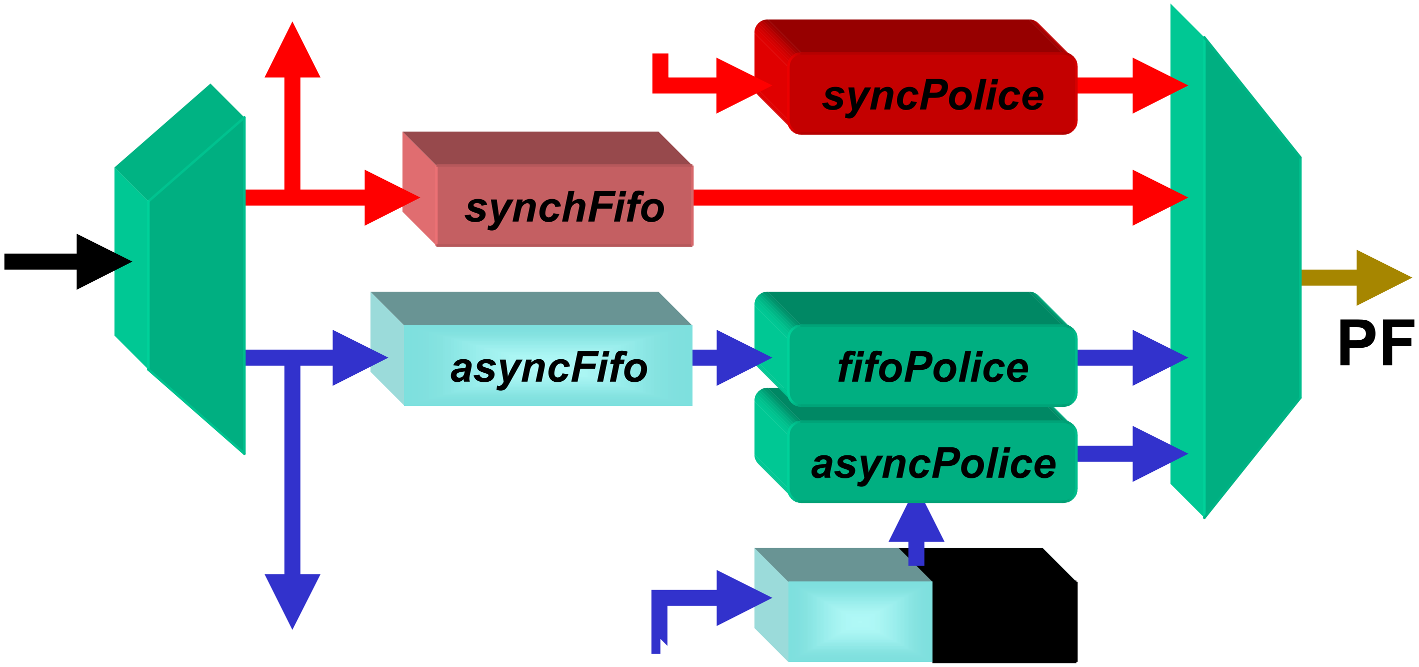

- Consumed asyncSend => inform others (asyncSend is higher level queue)

Lara Networks

Fairness assistance

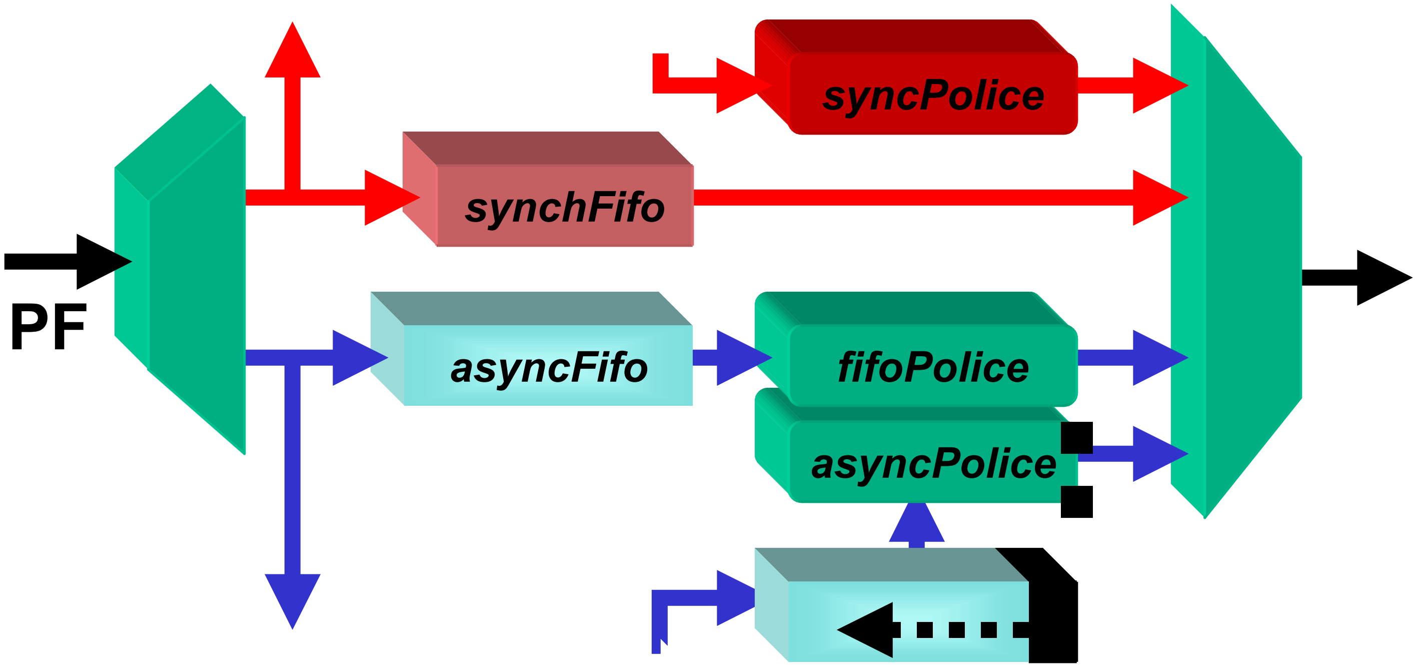

- Fair assistance => asyncSend filling
- Throttle nonprovisioned insertions

Token based fairness

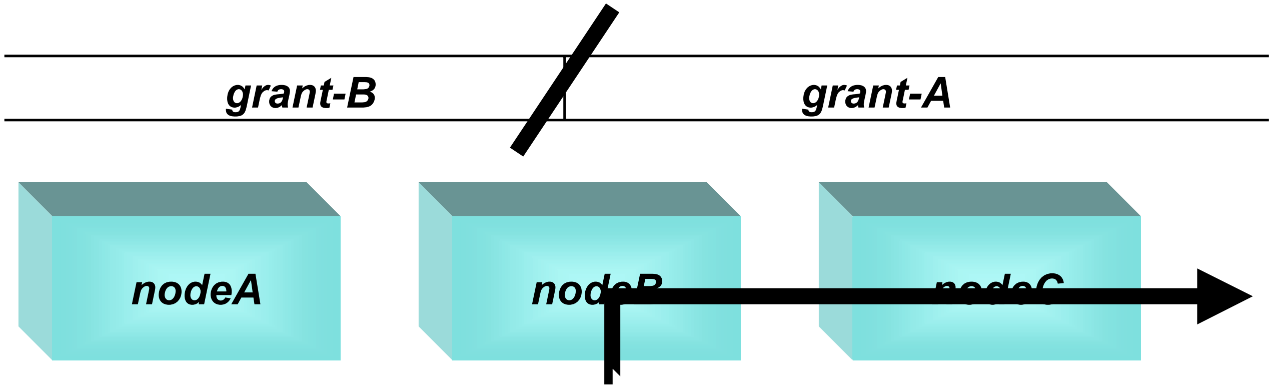

- Receive the grant wavefront
- Hold the wavefront while sending
- Fairness weighting is higher-level issue

Opposing arbitration

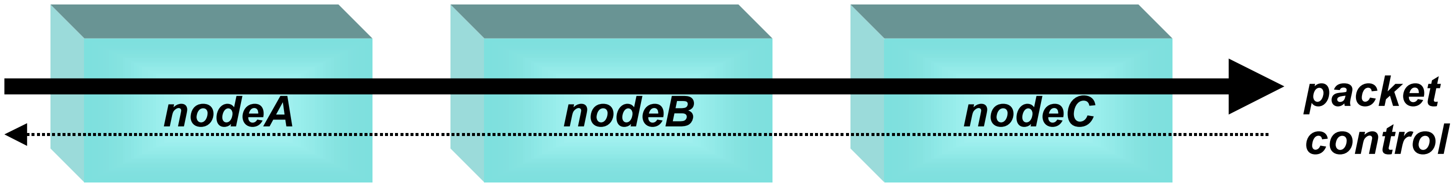

- Data packets flow in one direction
- Arbitration control flows in the other*

Lara Networks

Arbitration notes

- Dual levels
  - Synchronous, pre-emptive low latency
  - Asynchronous, negotiated and residual
- Jumbo frames
  - Affect asynchronous latencies
  - NO IMPACT on synchronous latency
- Idle symbols
  - Lowest latency for control
  - No catch-22
    (arbitrate for arbitration control)

For the future

- Heterogeneous bandwidths
  - Incremental growth useful (inevitable?)
  - QOS throughout the “ring”
- Concrete written proposals
  - Slide wars provide useful background
  - The devil is in the details...
  - Improves simulation credibility
  - “Open system” code-quality analogy
- Cut-through vs store-and-forward
  - Either should be allowed
  - Light-load latency DOES matter

Cut-through CRCs

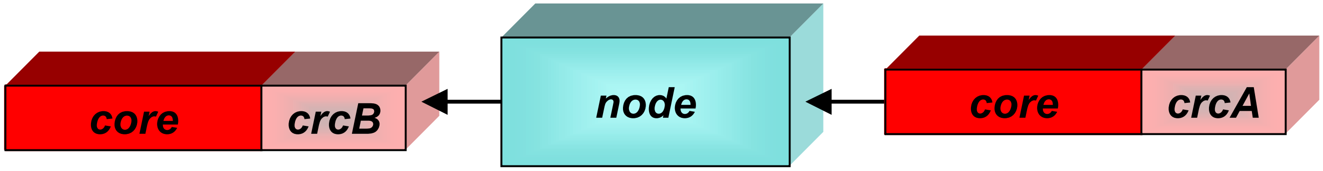

- Corrupted packet remains corrupted
- Error logged when first detected
- if (crcA!=crc&&crcA!=crc^STOMP) {

```
if (crcA!=crc&&crcA!=crc^STOMP) {
    errorCount+= 1;
    crcB= crc^STOMP;
}
```

Lara Networks