DOE High Performance Networks
Planning Workshop

Network Research: What is possible in the near-term, mid-term, long-term...

Mari Maeda
CISE NSF
August 13-15, 2002

Network Research and High Performance Network (HPN)

RESEARCH

Optical Networking (circuits)

Optical Networking (bursts)

QOS and Service Classes

High-Performance Protocols

Middleware Architecture & Tools

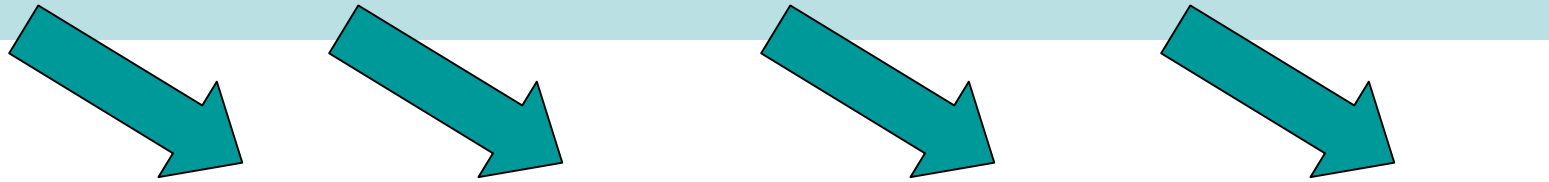

2002 2004 2006 2008 2010

High Performance Networking Infrastructure
(new generations of HPN enabled by net research)

High-Performance Networks TODAY (2002)

End-to-End Packet Loss and Throughput:

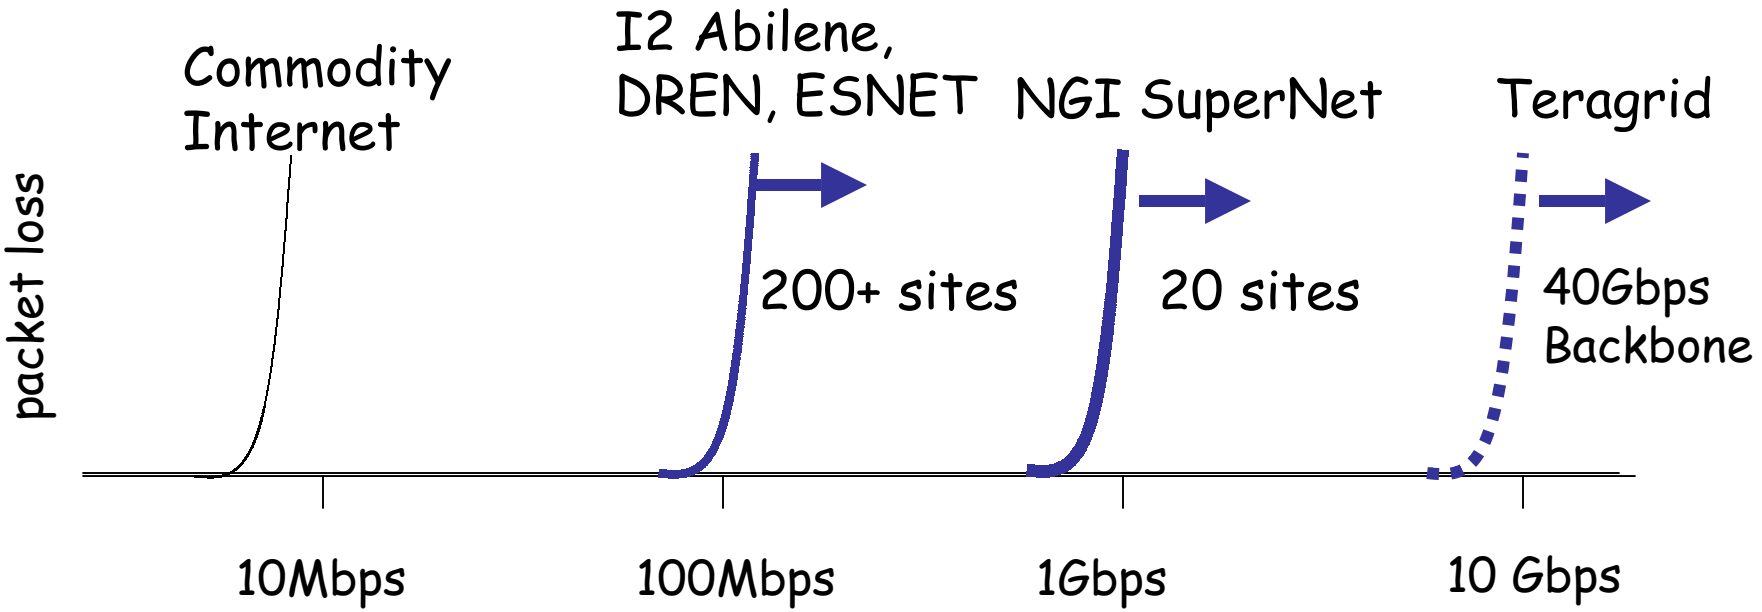

- While sites with multi-100's Mbps access is growing,
- Very FEW and ISOLATED sets of users with high-performance (Gbps) throughput

Challenge: Scaling High Performance Networks

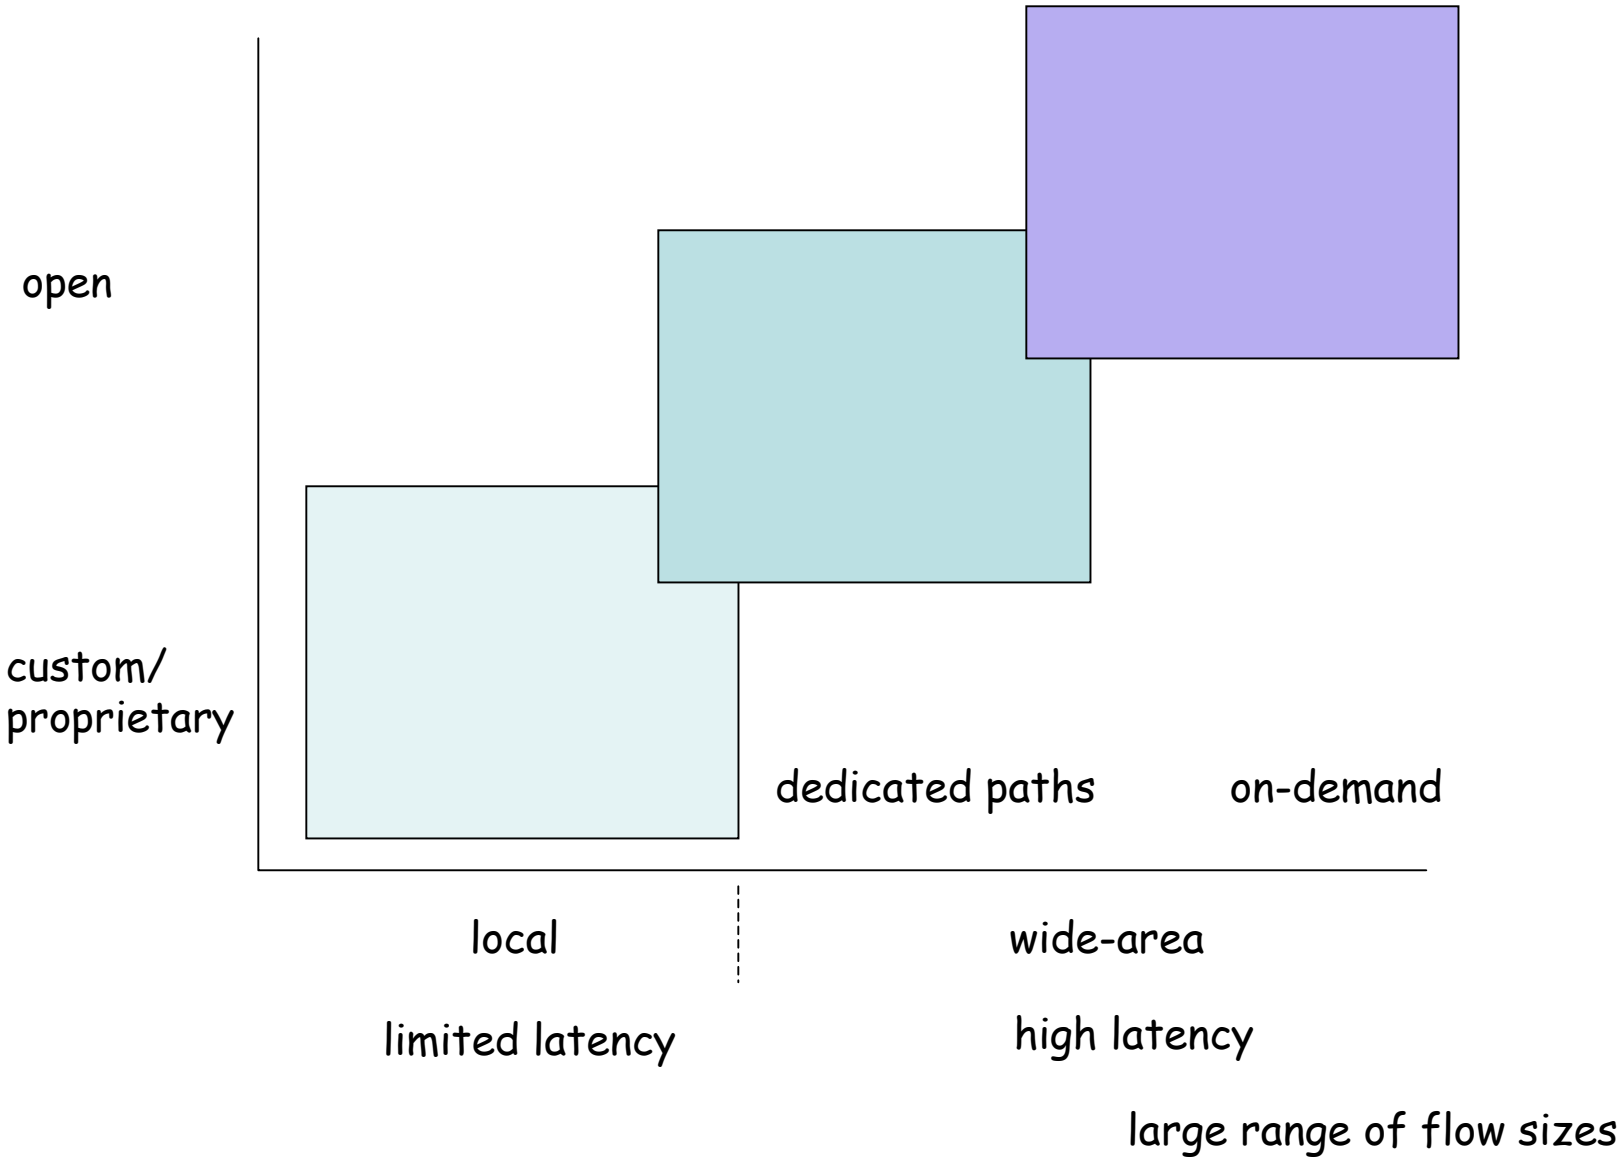

Limitations

High Performance Infrastructure needs scale:

- Push TCP-IP to Gbps, 10 Gbps ...

- Deployed circuit-based optical networking architecture is not scalable
(today's optical networking architecture, even with GMPLS, is coarse-grained and cannot cost-effectively simultaneously support very many users)

Issues (near- to mid-term)
Scaling TCP-IP to Gbps

- Today, limited support for large MTU's
- Increasing MTU's and automatic tuning for large flows may be the first step.
- Need understanding of interactions between MTU size, slow start, TCP flow control etc.
- Stability problem of TCP at large window sizes necessary in high latency, high capacity environment.
- Exploration of new queue management algorithms - via simulation and experimental testing

Optical Networking

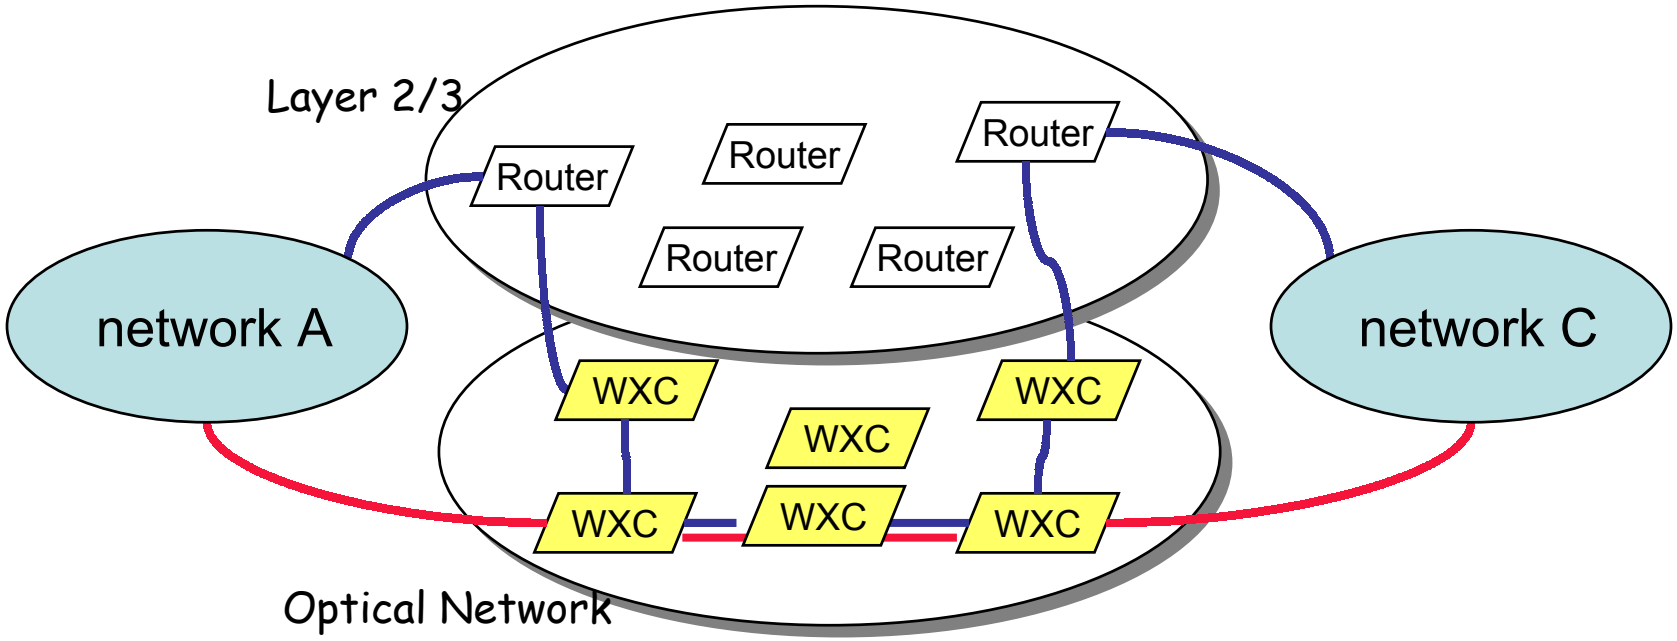

Features of greatest interest to service providers:

- network protection/restoration
- fast provisioning of services (from weeks/months to hrs/days)
- traffic engineering

Feature that's of interest to High Performance Networks:

- bottleneck bypass
- possibly, bandwidth-on-demand

Optical Networking

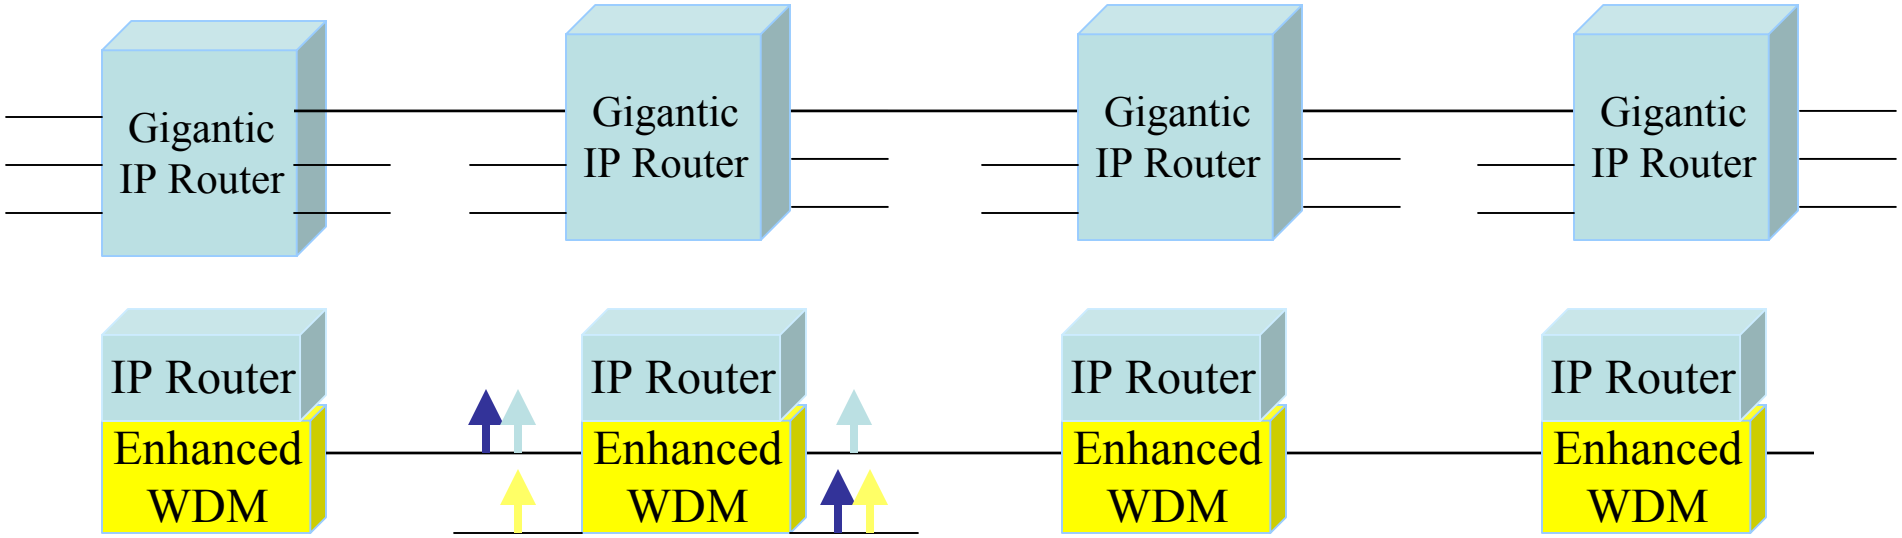

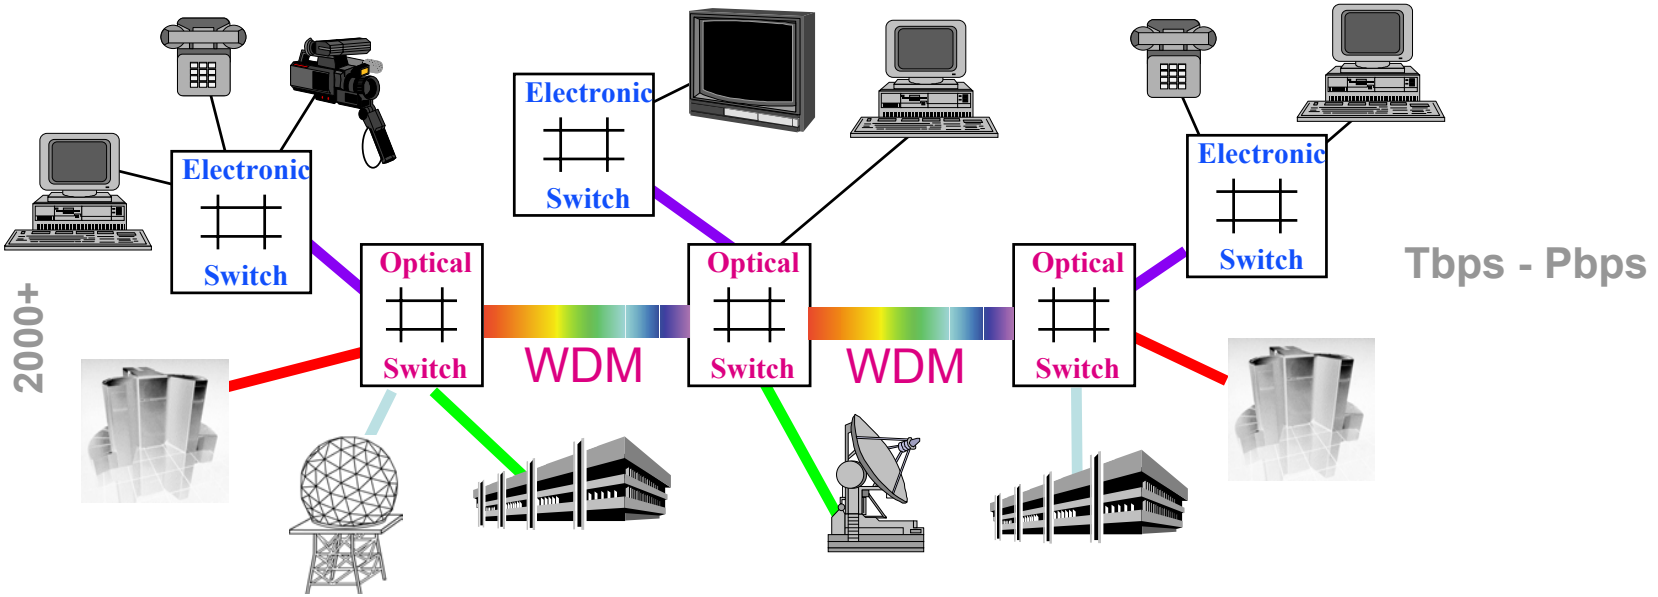

Optical Path Configuration

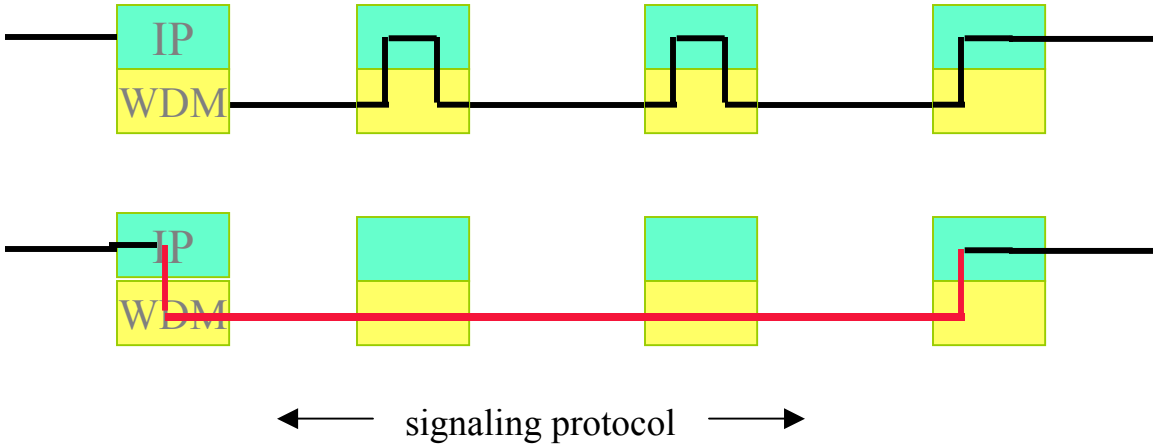

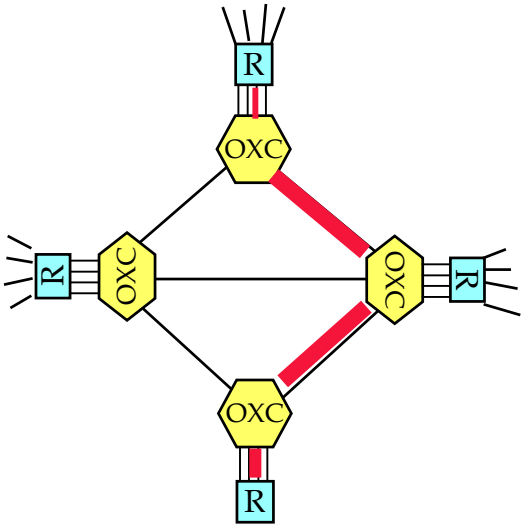

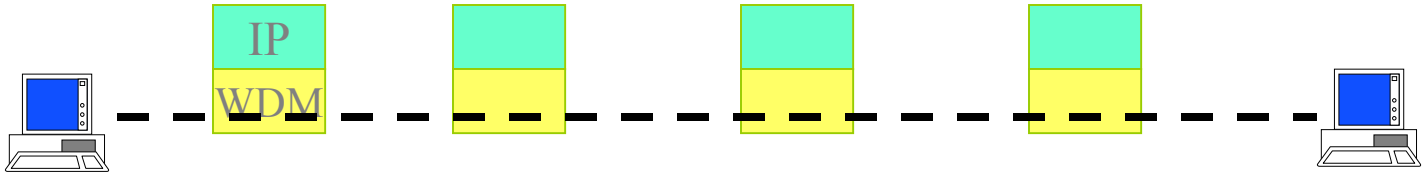

dynamic optical switching to off-load electronic routers.

ONRAMP Experimental Network

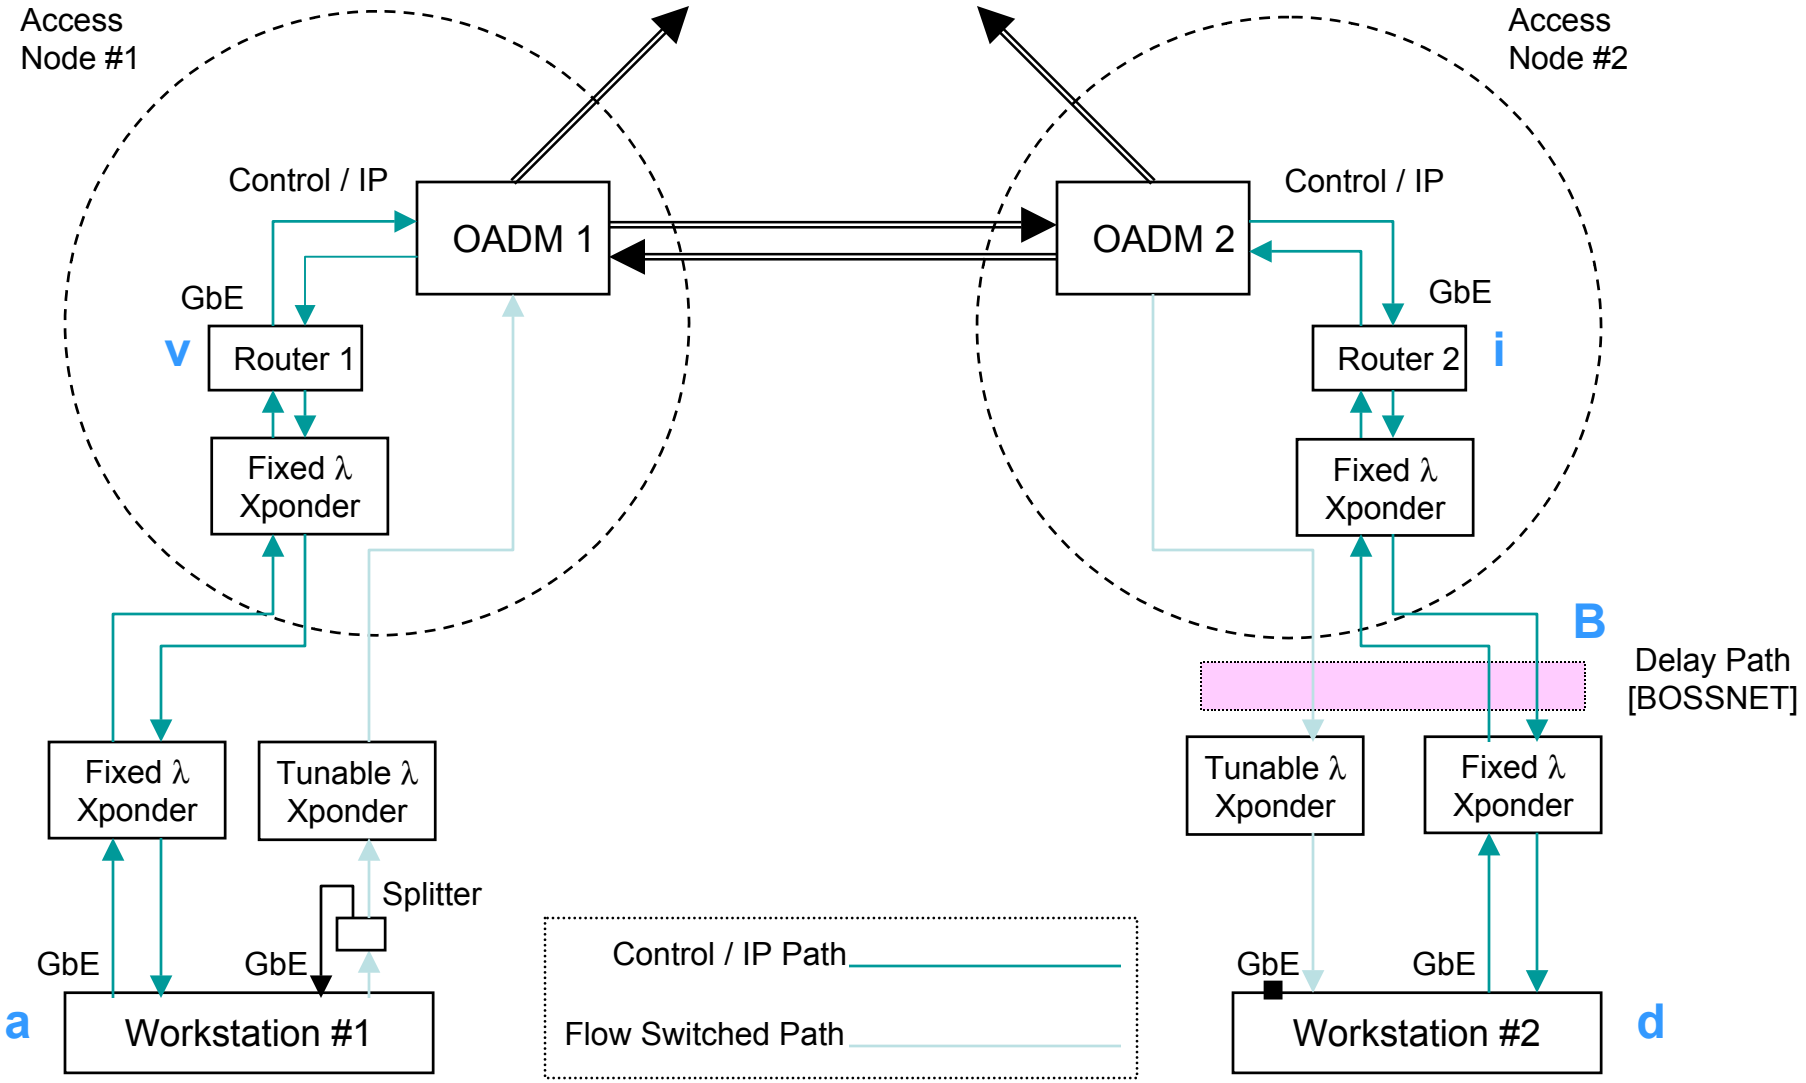

MIT Lincoln Laboratory

Example:
High-Performance Research Network
End-To-End Performance and Application Experiments

- ATDNET-BOSSNET-GLOWNET Testbed
- ~1000 miles roundtrip; prototype networking system hardware
- Optical spectra (multiple wavlengths) for optical transmission experiments
- A wavelength dedicated for Gps+ TCP-IP experiments
- Real-Time Gbps Applications Experiments

Flow Switch to higher rate path

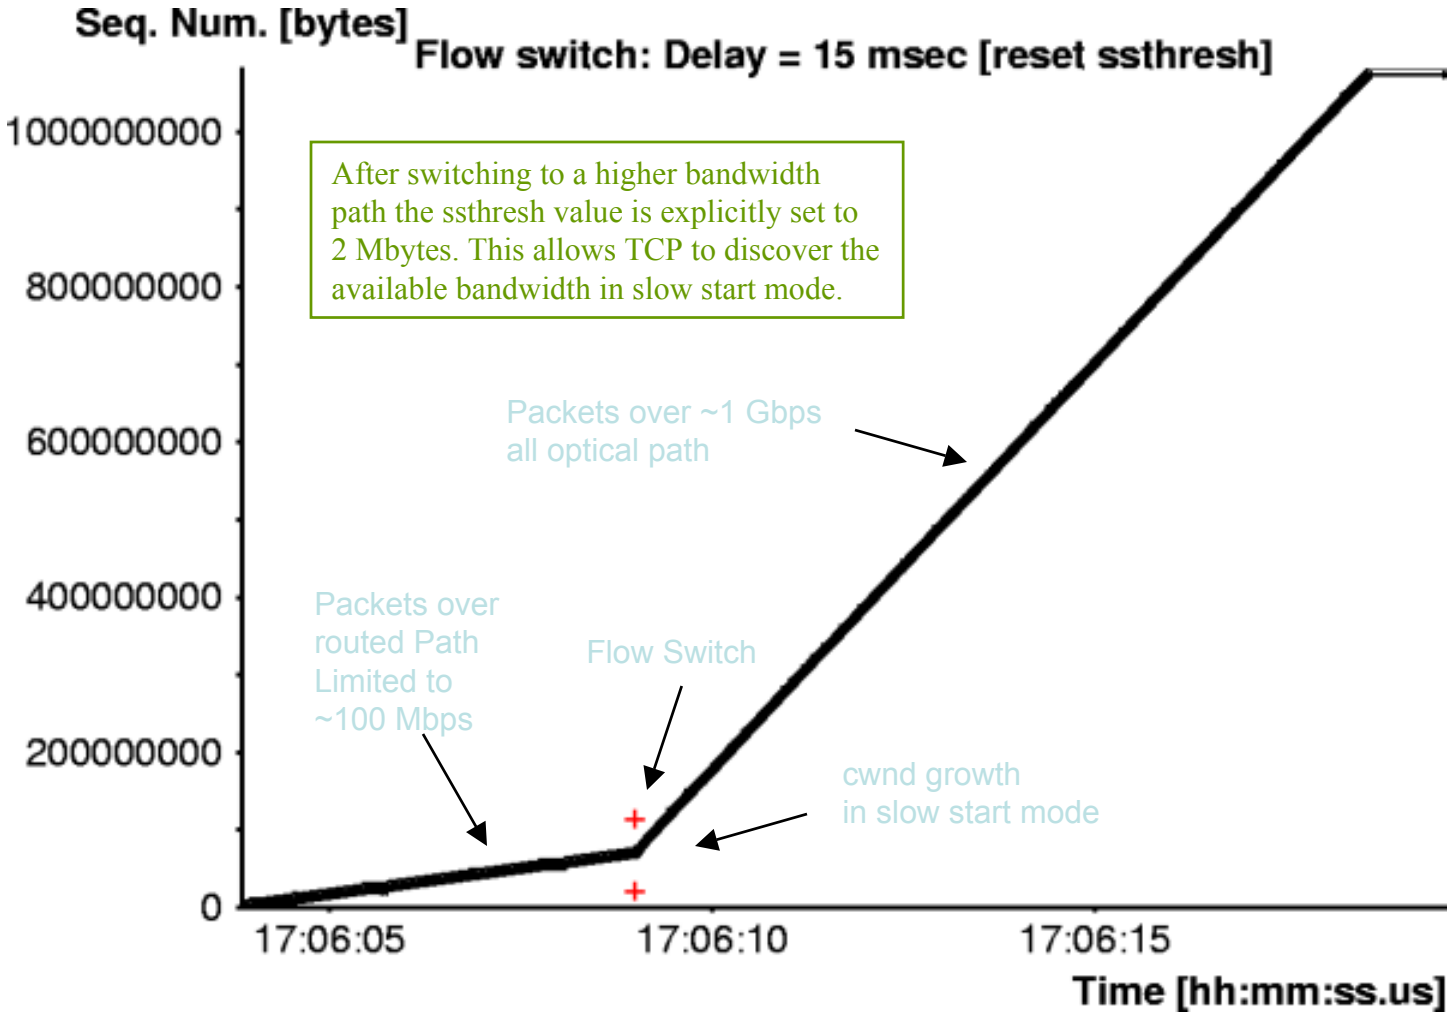

Fig. 8a

ssthresh reset to 2 Mbyte

Statistical Multiplexing in the Optical Layer

	holding time
Reconfigurable Opt. Networking (Circuit Switching)	**days, months**
Optical Burst Switching	**>μsec**

assuming 2.5 Gbps end to end

1 sec :	**300 MB**	**(DOQ)**
10 msec :	**3 MB**	**(MP3)**
30 μsec :	**10KB**	**(web page)**

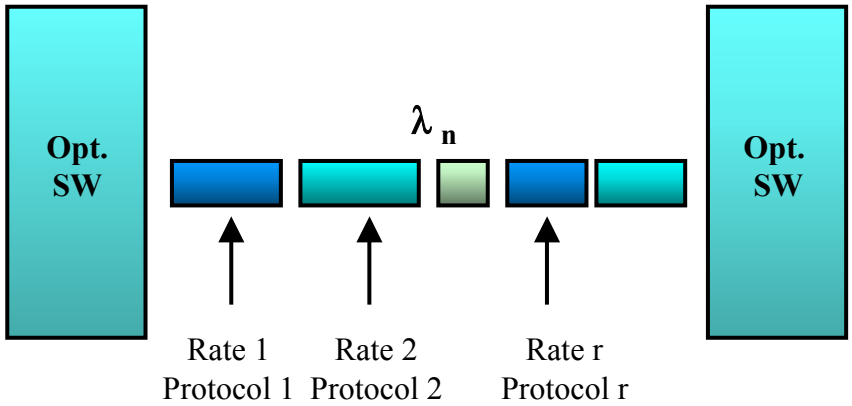

Goal:
Bit rate and protocol agile bursts

research challenges

- speed and size of optical switches
- gain stabilization of amplifiers or optical "idle" to minimize transients
- fast clock recovery at variable bit rates
- control plane (one-way, two-way reservation, header)
- buffering, contention resolution, flow control...

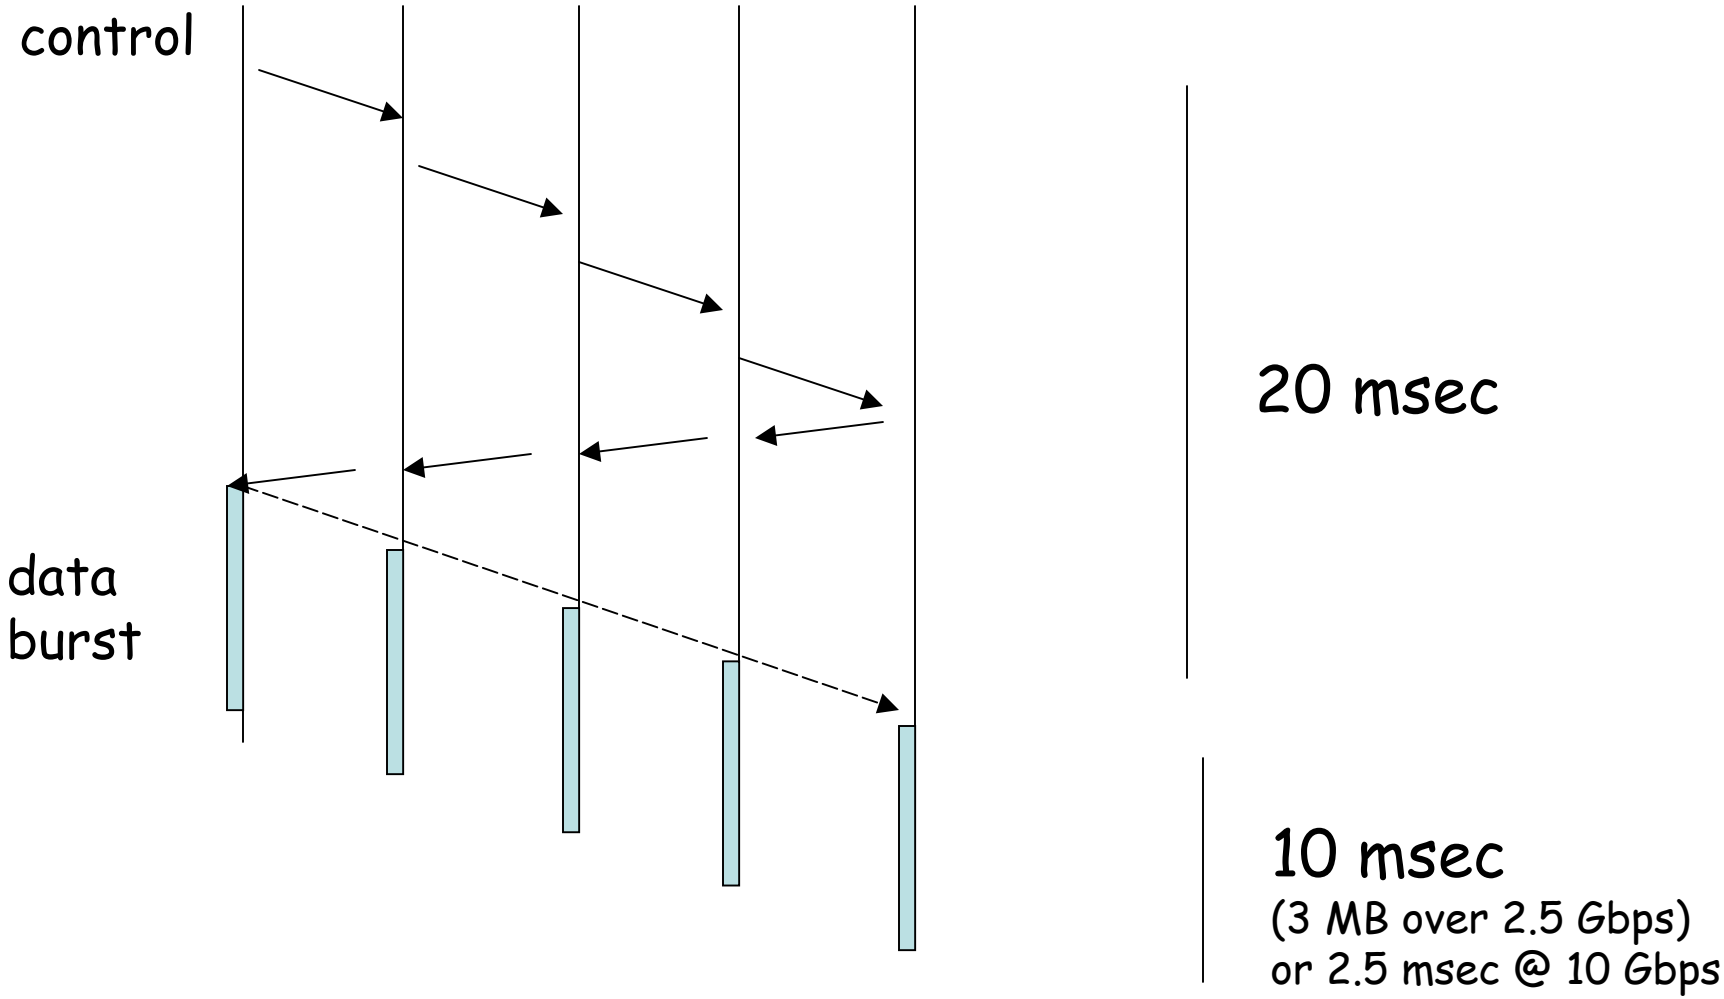
control
data
burst
20 msec
10 msec
(3 MB over 2.5 Gbps)
or 2.5 msec @ 10 Gbps

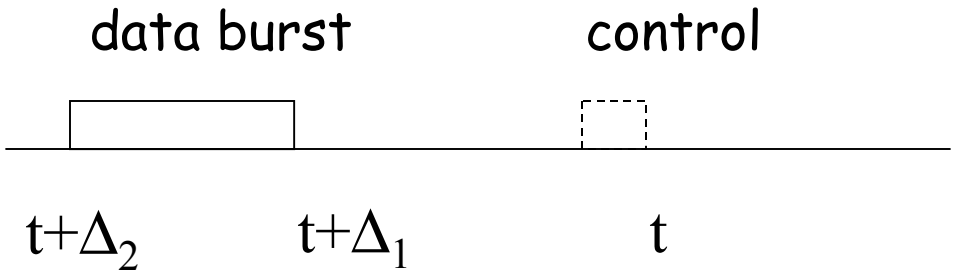

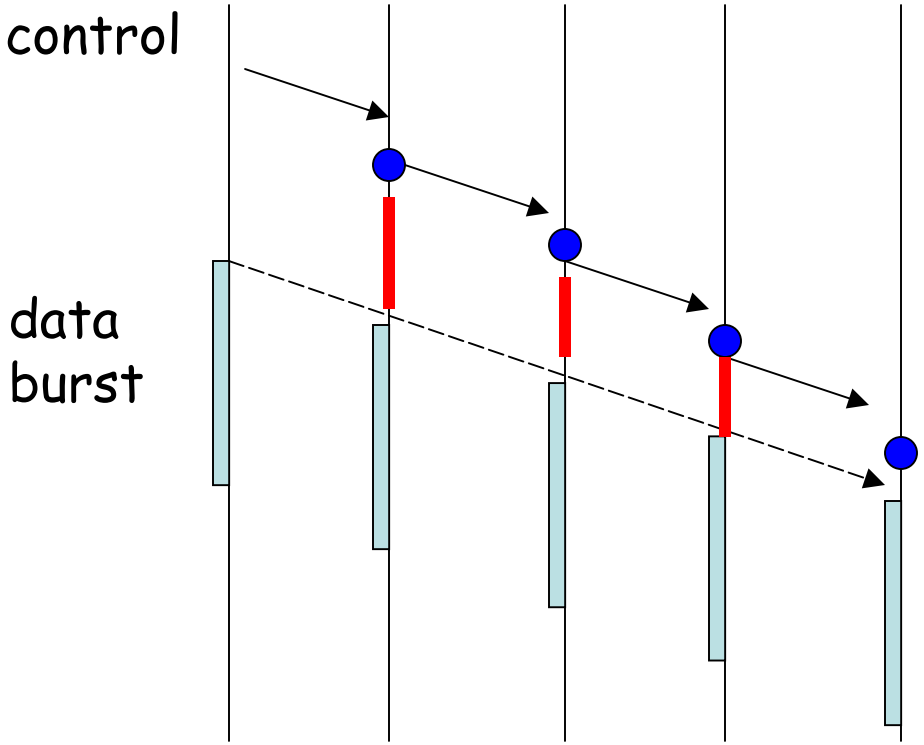

● crossconnect activation

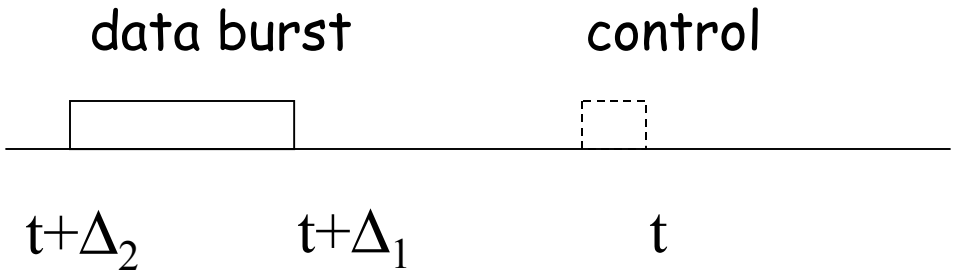

- estimated burst arrival time
- burst length

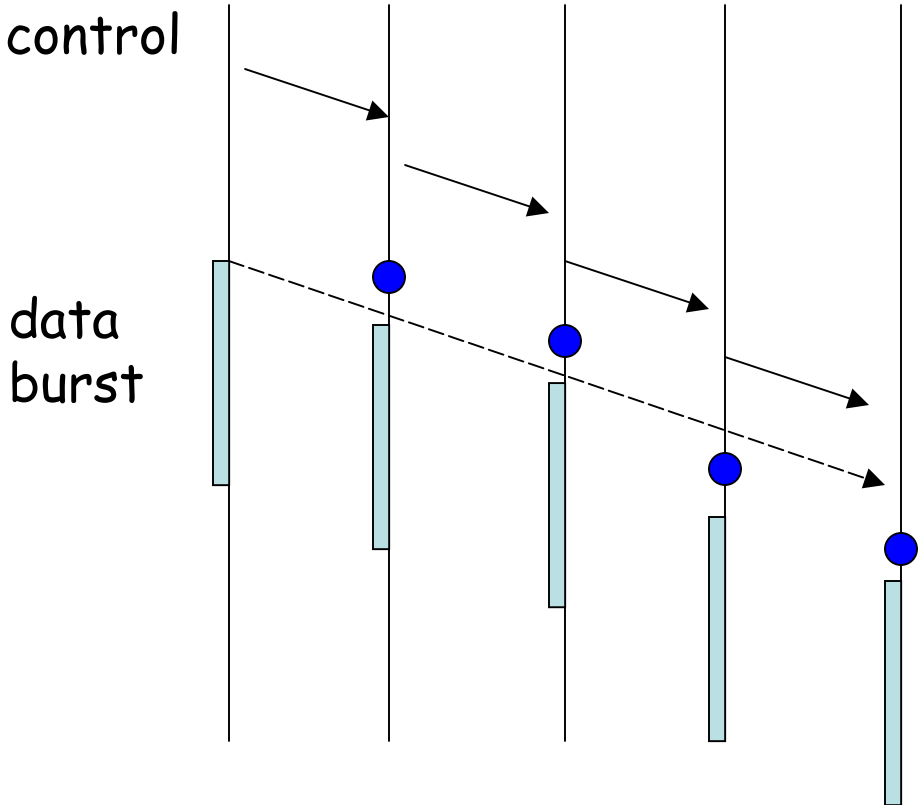

● crossconnect activation

J. Turner, Washington University OBS Project

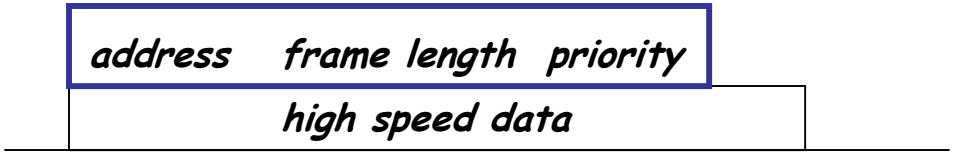

data burst & label (header)

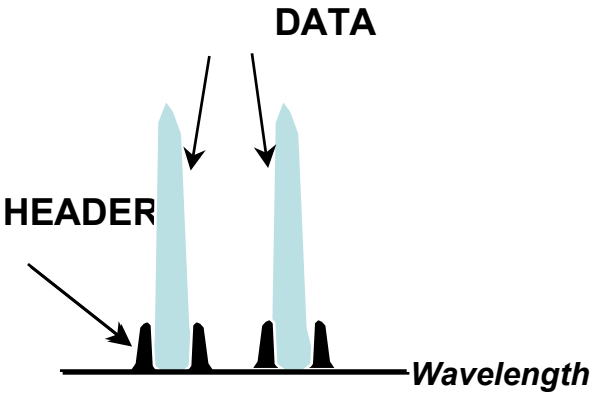

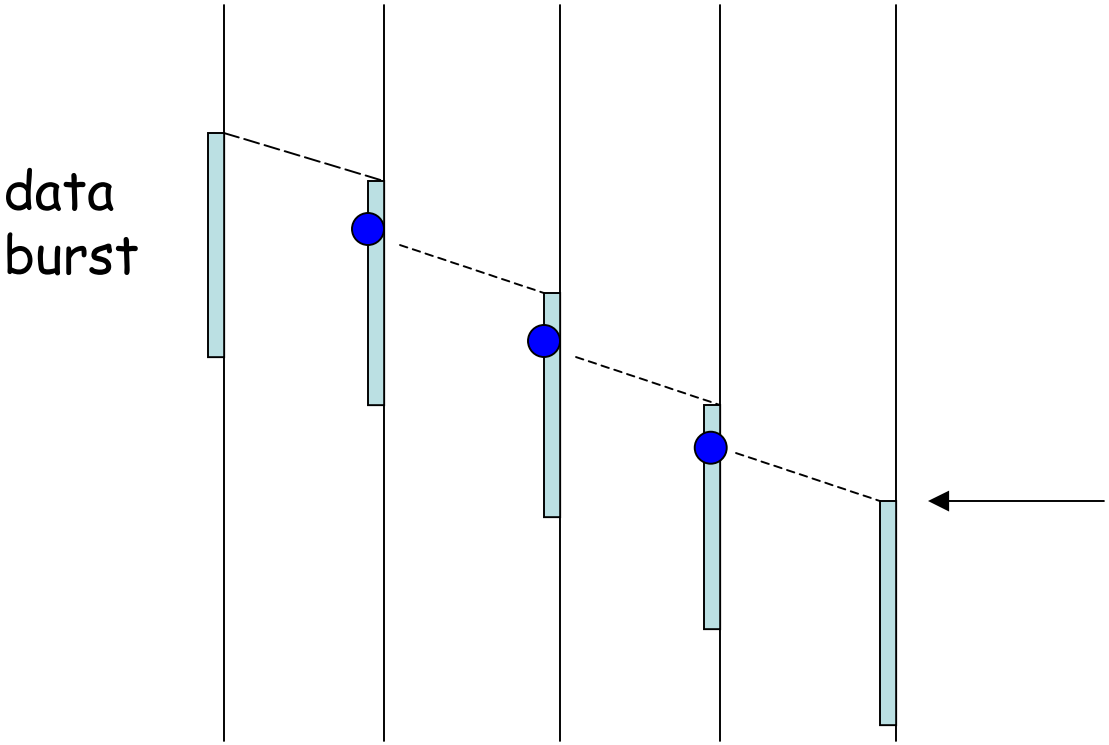

burst arrival time

Future Challenges
Real-Time Applications

Today (limited bw):
Data collection followed by local data reduction

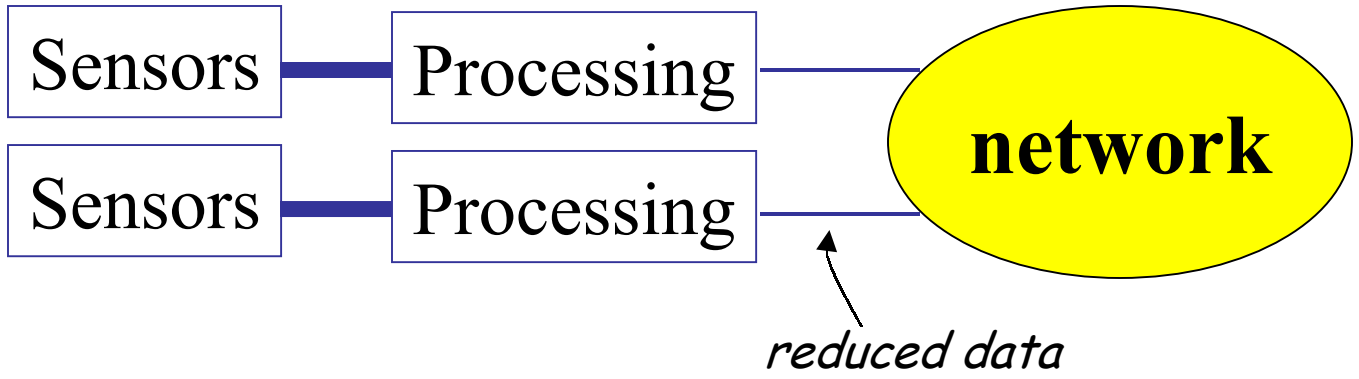

Challenge:
Networked collaborative/coherent processing of raw data enables new capabilities

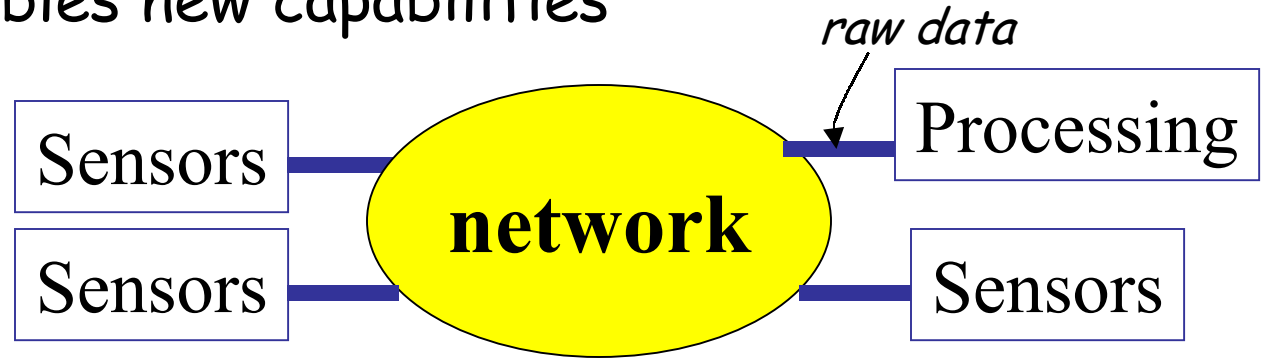

VLBI Science

ASTRONOMY

- Highest resolution technique available to astronomers – tens of microarcseconds
- Allows detailed studies of the most distant objects

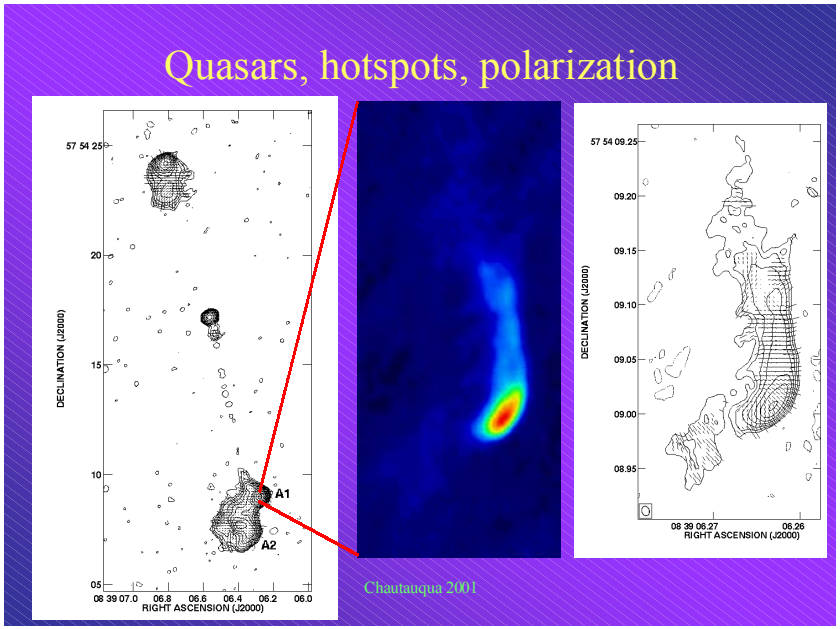

GEODESY

- Highest precision (few mm) technique available for global tectonic measurements
- Highest spatial and time resolution of Earth's motion in space for the study of Earth's interior
  - •Earth-rotation measurements important for military/civilian navigation
  - •Fundamental calibration for GPS constellation within Celestial Ref Frame

Plate-tectonic motions from VLBI measurements

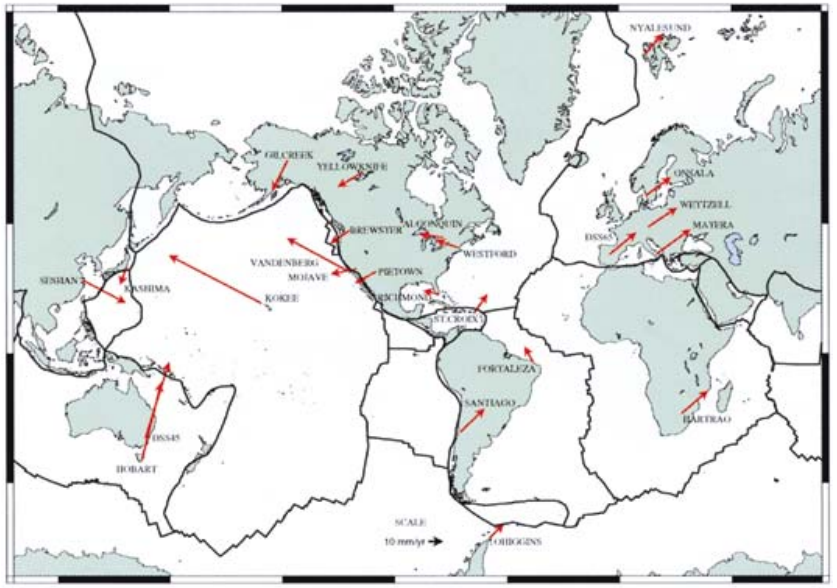

Alan Whitney, MIT

Scientific Advantages of e-VLBI

- Bandwidth growth potential for higher sensitivity
  - VLBI sensitivity (SNR) proportional to square root of Bandwidth resulting in a large increase in number of observable objects
    (only alternative is bigger antennas – hugely expensive)
  - e-VLBI bandwidth potential growth far exceeds recording capability
    (practical recordable data rate limited to ~1 Gbps)
- Rapid processing turnaround
  - Astronomy
    - Ability to study transient phenomena with feedback to steer observations
  - Geodesy
    - Higher-precision measurements for geophysical investigations
    - Better Earth-orientation predictions, particularly UT1, important for military and civilian navigation

- Given abundant capacity and latencies limited by "c", what new capabilities are enabled?
- What network research will enable new generations of HPNs?
- Given that capacities and latencies are non-ideal, what technologies and tools may optimize the use of our under-utilized resources?
  - BW
  - CPU: large scale public computing
    - general purpose framework and infrastructure for public-resource computing
    - not necessarily over HPN
    - SETI@home user base 3.5 million people worldwide
    - new open source architecture and open-source sw to allow multiple independent projects to share a single participant base

Future Research Networks

Operational Research Network

Experimental Infrastructure Network

driven by applications

push leading edges of performance

Research Network Testbed

enables new technology research, development and prototyping

Summary

- Scalable HPN infrastructure critical for support of scientific applications and national security applications with growing data volumes.

- Research in networking (physical layer, high-perf protocols, middleware, architecture, tools) defines new generations of high-performance networks

- Near-term emphasis on: large MTU's, improvements to TCP-IP to handle multi-Gbps, combined with optical circuits

- Investigation beyond " TCP-IP " needed for long-term HPN evolution.