After the bandwidth glut

Farzad Safaei

April 6, 2005

Smart Internet Technology CRC
Centre for Emerging Networks and Applications
University of Wollongong

Cross-over points we left behind (1)

Voice vs Data Traffic

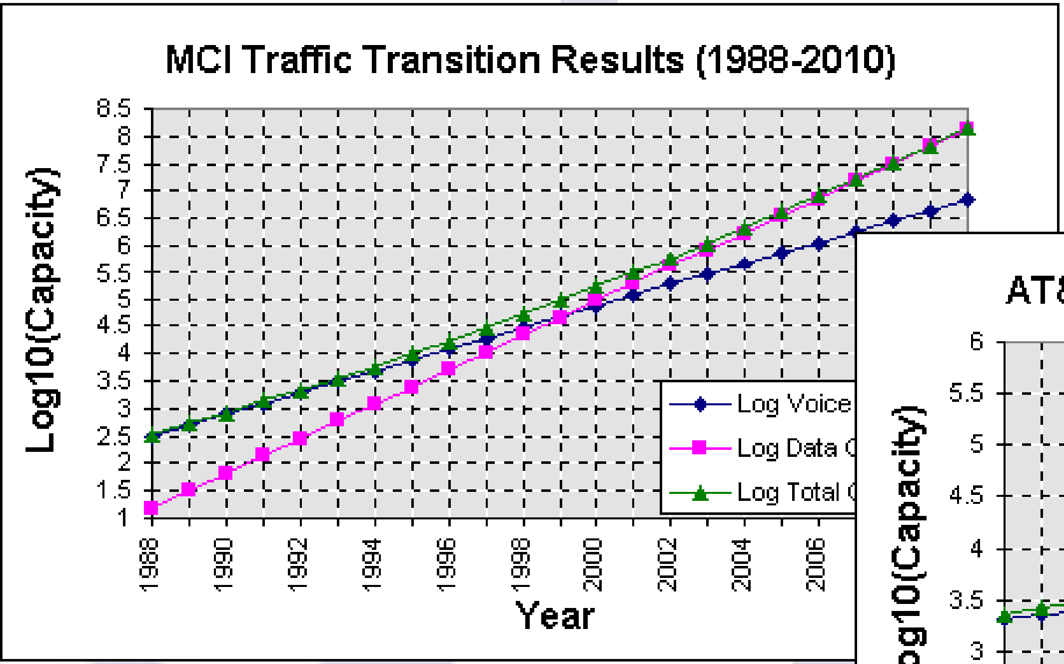

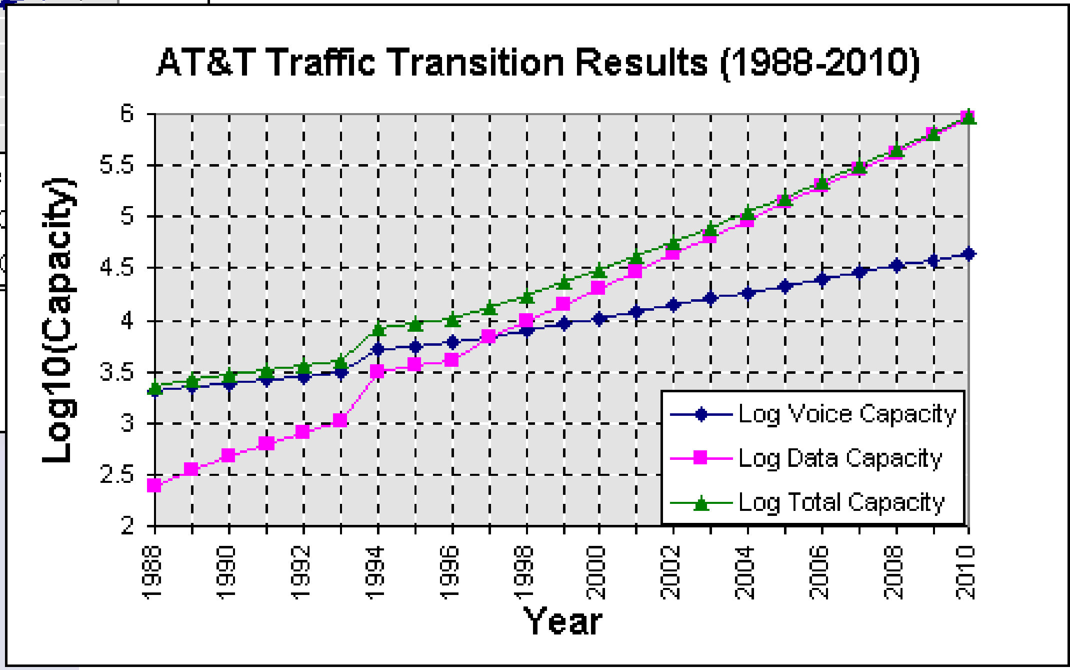

Voice & Data Capacity in Mbps
Source: Mutooni & Tennenhouse, MIT, Jan 1998

Cross-over points we left behind (2)

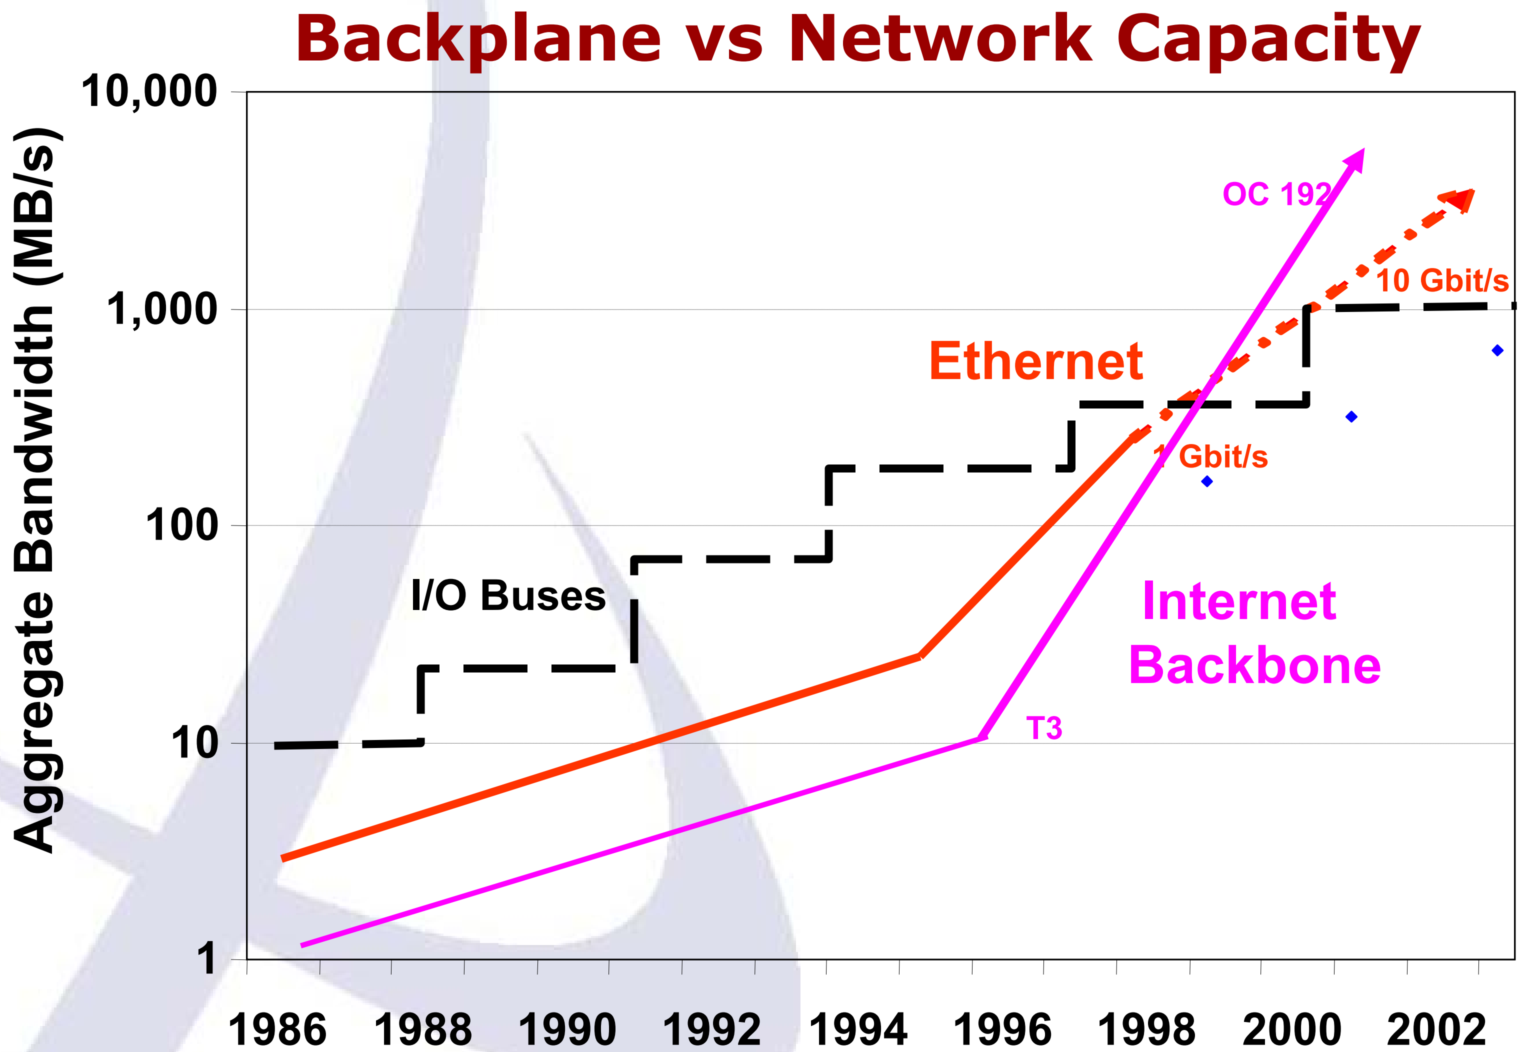

Source: Openetlab, Nortel Networks, 2000

Cross-over points we left behind? (3)

Fixed vs Mobile Numbers

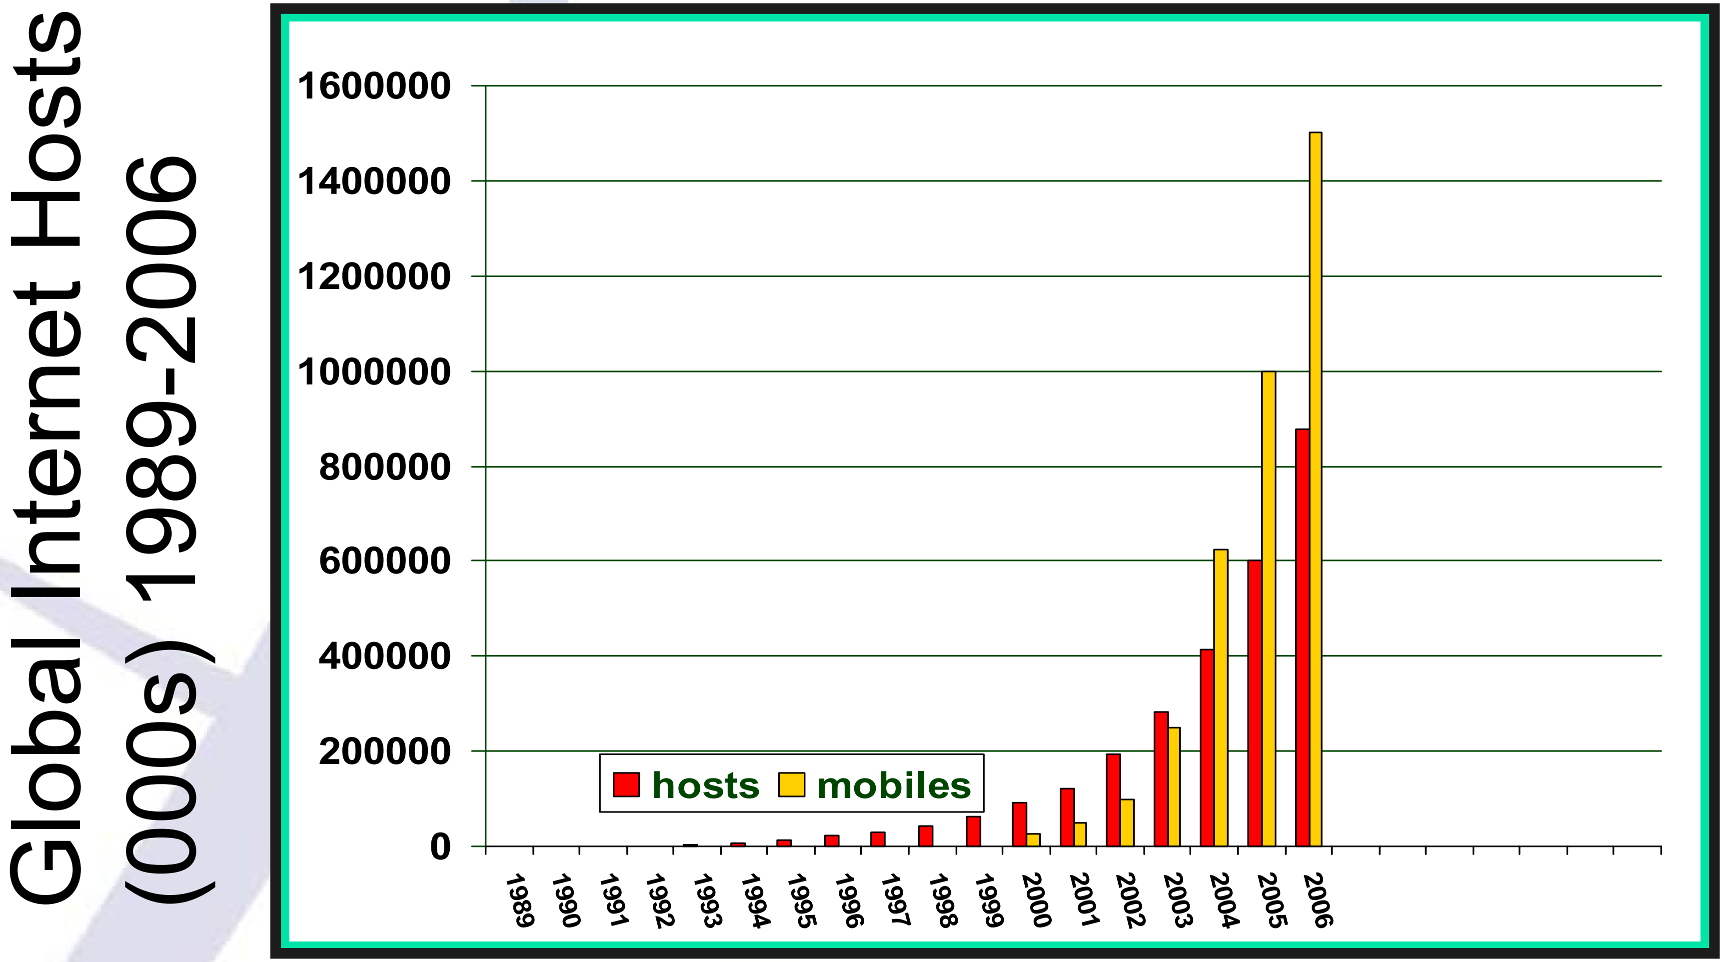

Source: Vint Cerf, MCI Worldcom, Jan 2000

Cross-over points ahead of us (1)

Peer-to-Peer vs Web Traffic

- At least for a while P2P traffic on some backbones exceeded web traffic
  - Measurement in European backbone showed around 60% of traffic being P2P related

Source: Ben Azzouna, N., Guillemin, F., “Impact of Peer-to-Peer Applications on Wide Area Network Traffic: An Experimental Approach”, Proceedings of IEEE Globecom 2004, Dallas, USA, 29 Nov - 3 Dec 2004.

Cross-over points ahead of us (2)

Access vs Core Bottleneck

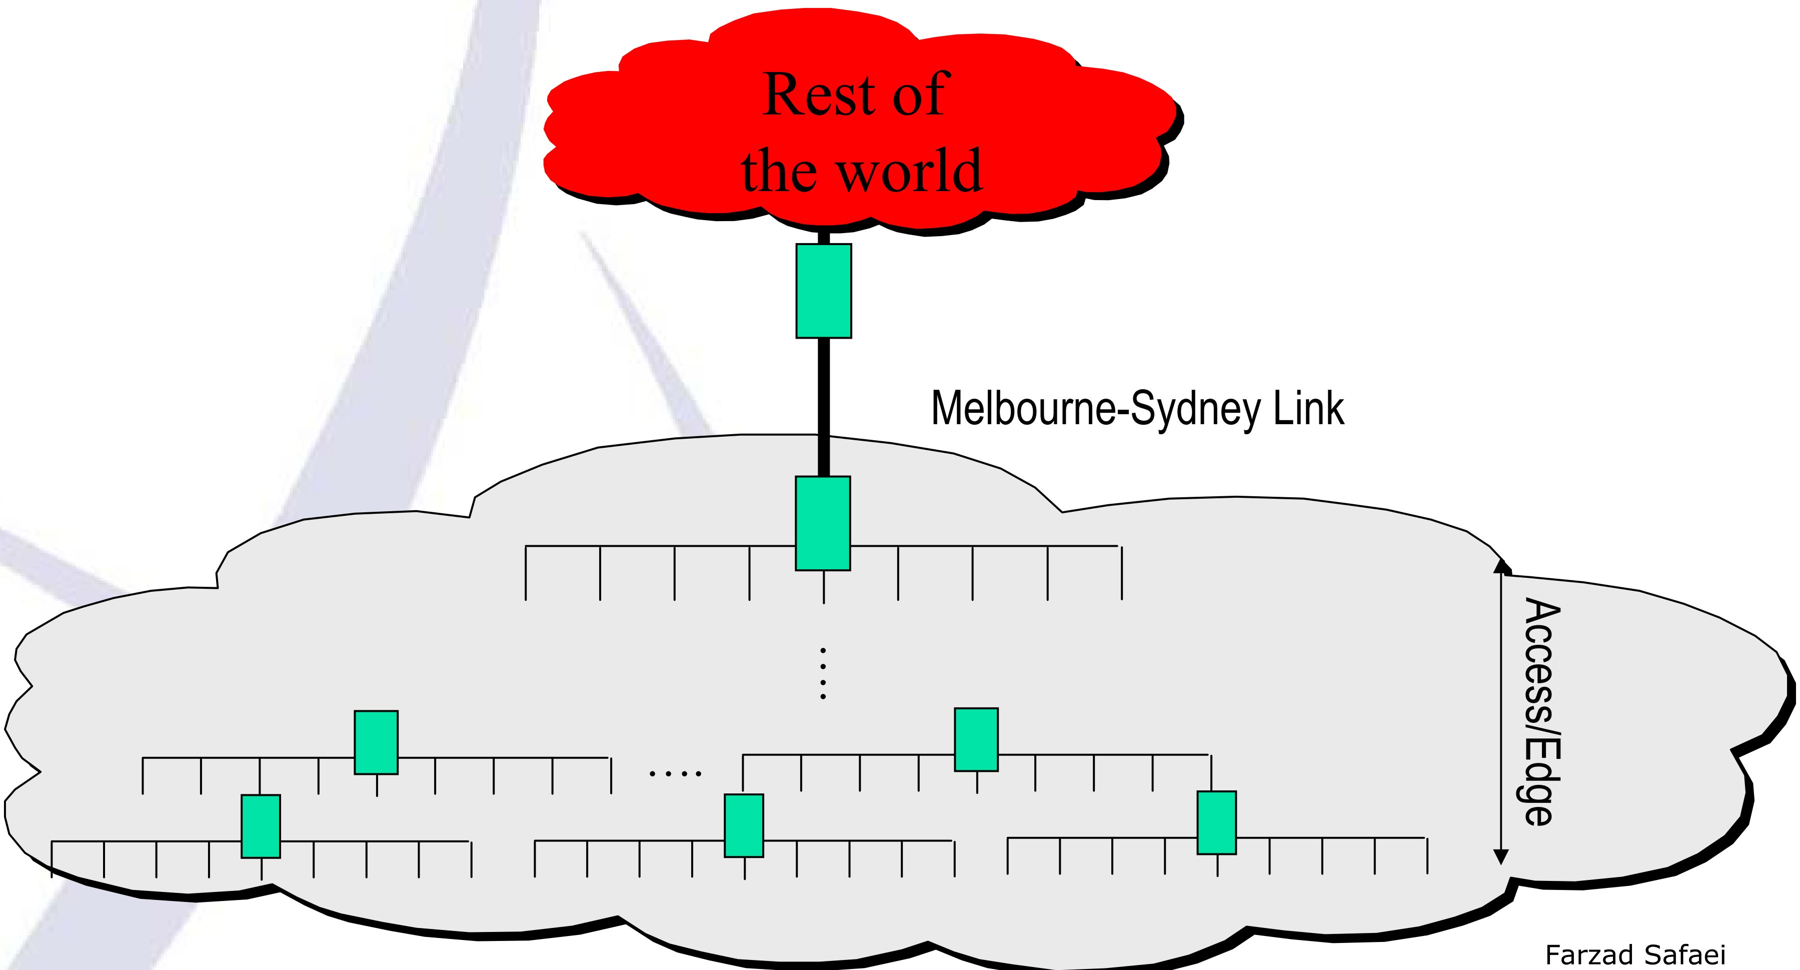

Economic constraints

- the 'scale of investment' will be different in different network segments
- the 'payback periods' will be different
- Vendor's constraints - 'quantum' of bandwidth that can be purchased will get larger

Demand as determined by the access capacity

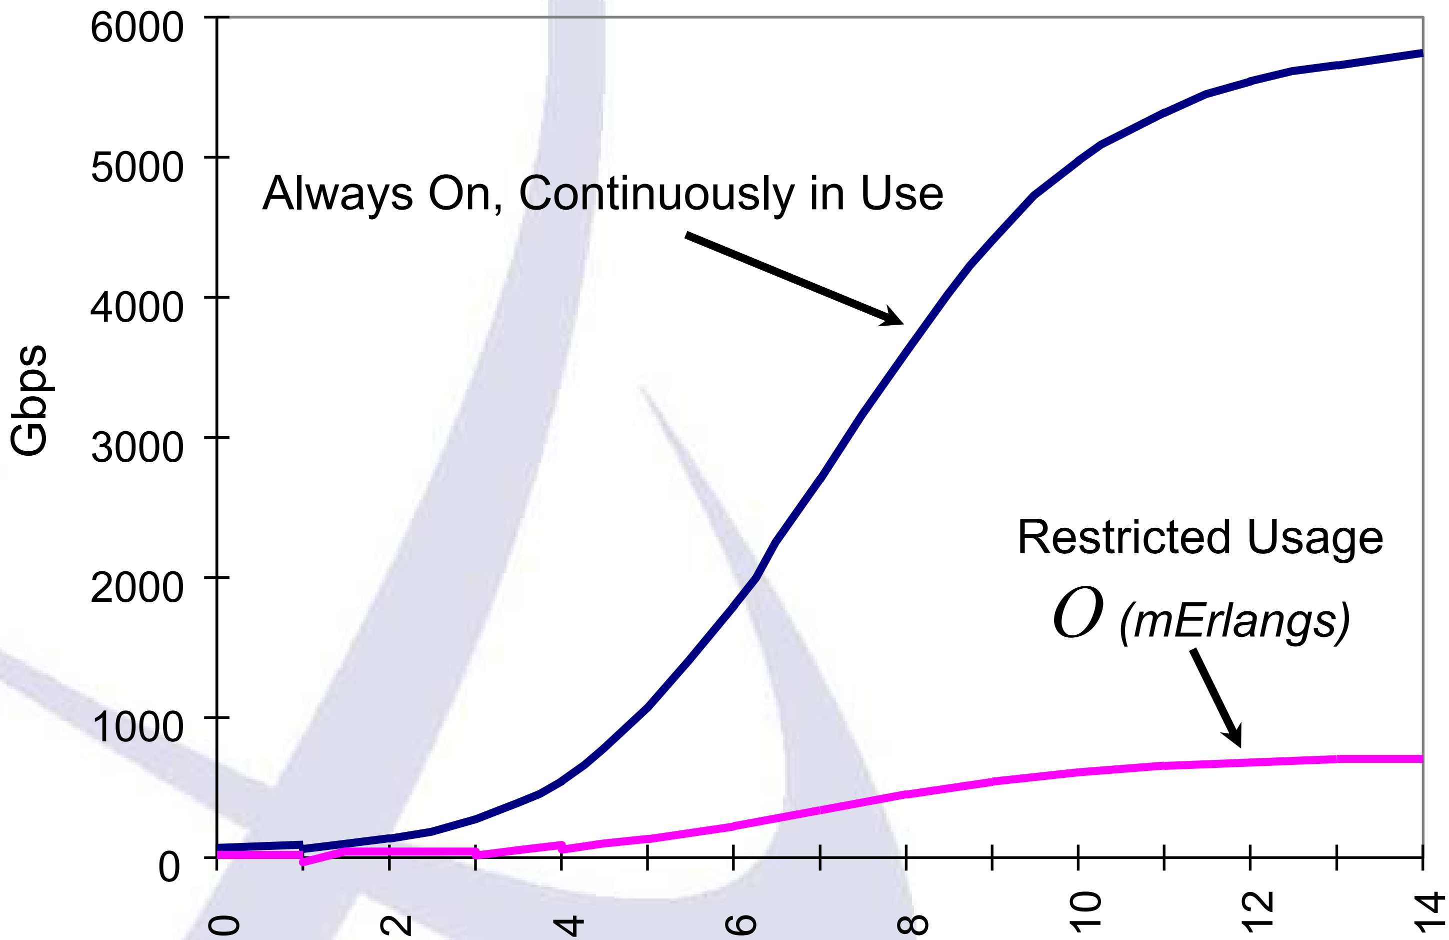

BB access at 65% in 10 yrs

BB saturation at 80%

BB average usage= 6 M

NB average usage= 56 k

Core/Access Bottleneck Crossover

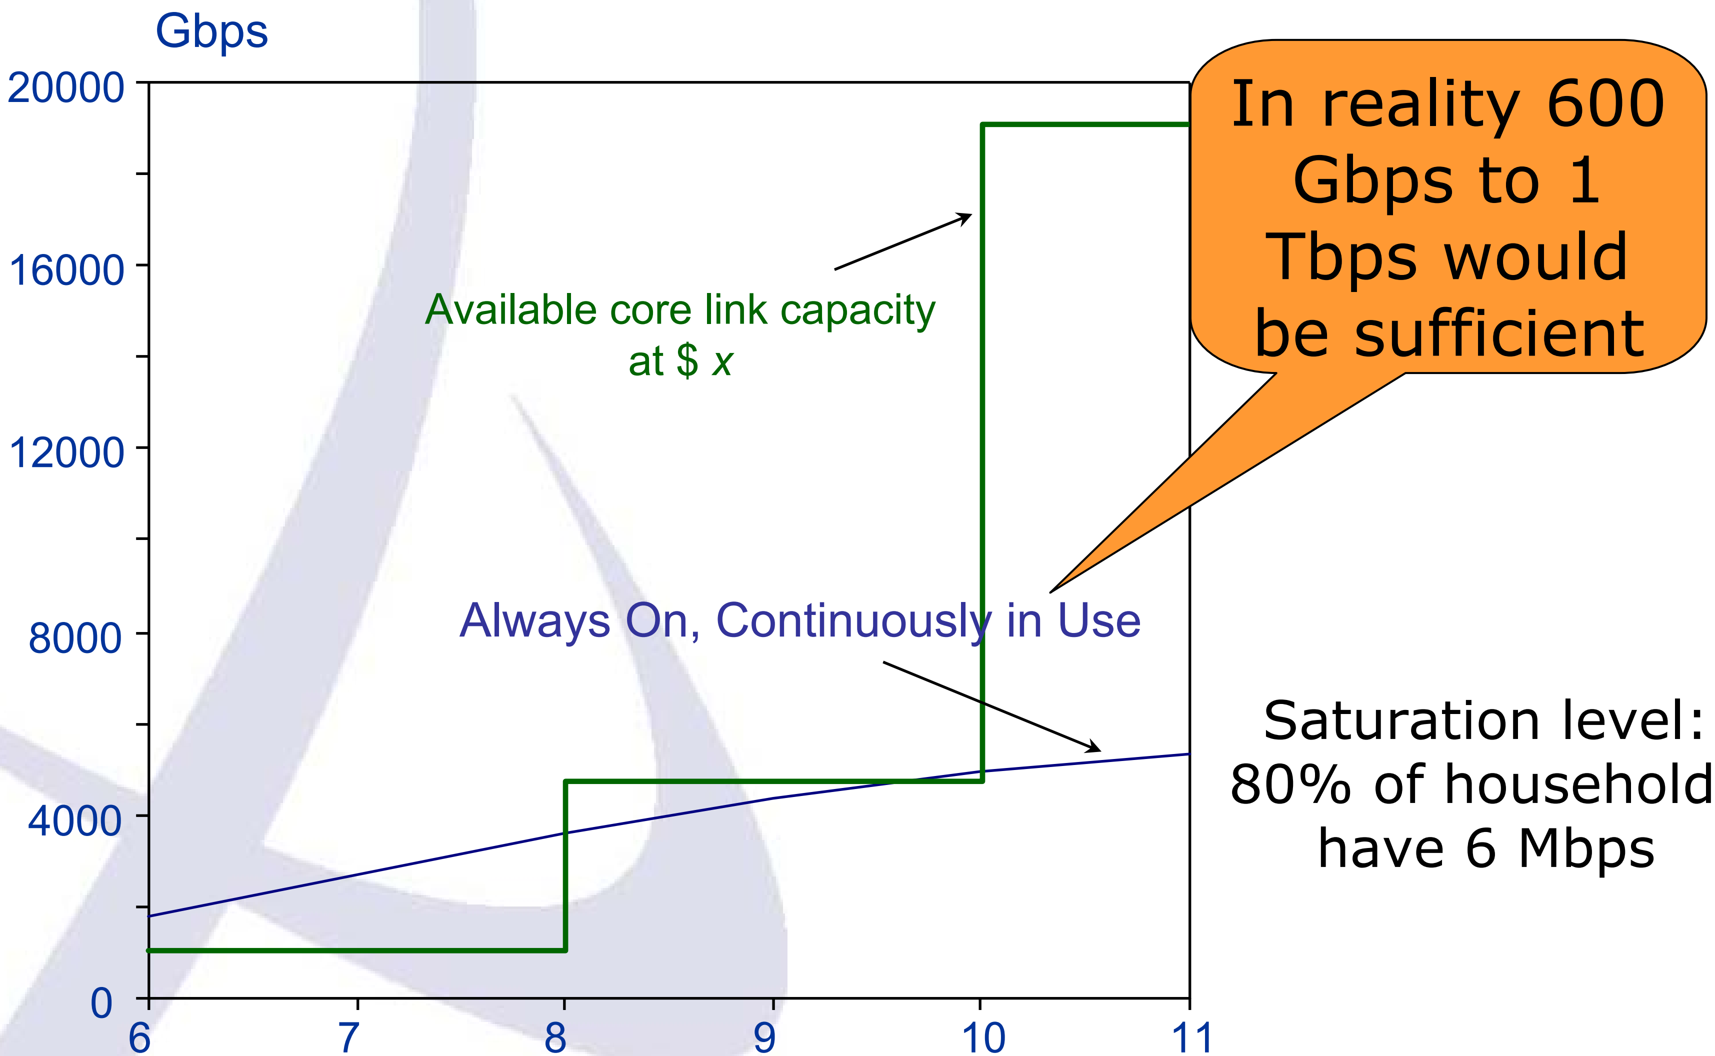

Saturation level: 80% of households have 6 Mbps

Cross-over points ahead of us (3)

Physiological Constraints

- Limits of human physiology for consumption of multimedia information
- Assuming immersive participatory applications could fund the access rollout, then based on some reasonable assumptions:

The Melbourne-Sydney link requires 935 Tbps

Cross-over points ahead of us (3)

Physiological Constraints

- With end-to-end optical
  - Assume no multiplexing gain
  - And circuit switched optical

The Melbourne-Sydney link requires 14 Pbps

Assumptions:

- Near abundance of bandwidth within the core
- Broadband access rollout will continue

Focus:

- What classes of applications will emerge?
- What are the new bottlenecks?
- What new features are needed?

Entertainment and Internet

- Currently a significant portion of Internet traffic is entertainment related
  - Primarily retrieval
- New entertainment applications are becoming very popular
  - Already the main driver for broadband uptake
  - Require capabilities which are very different

Participatory Applications

- A subset of interactive entertainment
  - Delivered over the Internet or Digital TV
- Interactivity involves immersion and virtual participation
- Example: Networked games
  - Estimated to have 230 M active players by 2009
- Other domains: collaboration, education, …

Multiple Inter-dependent Real-time Flows

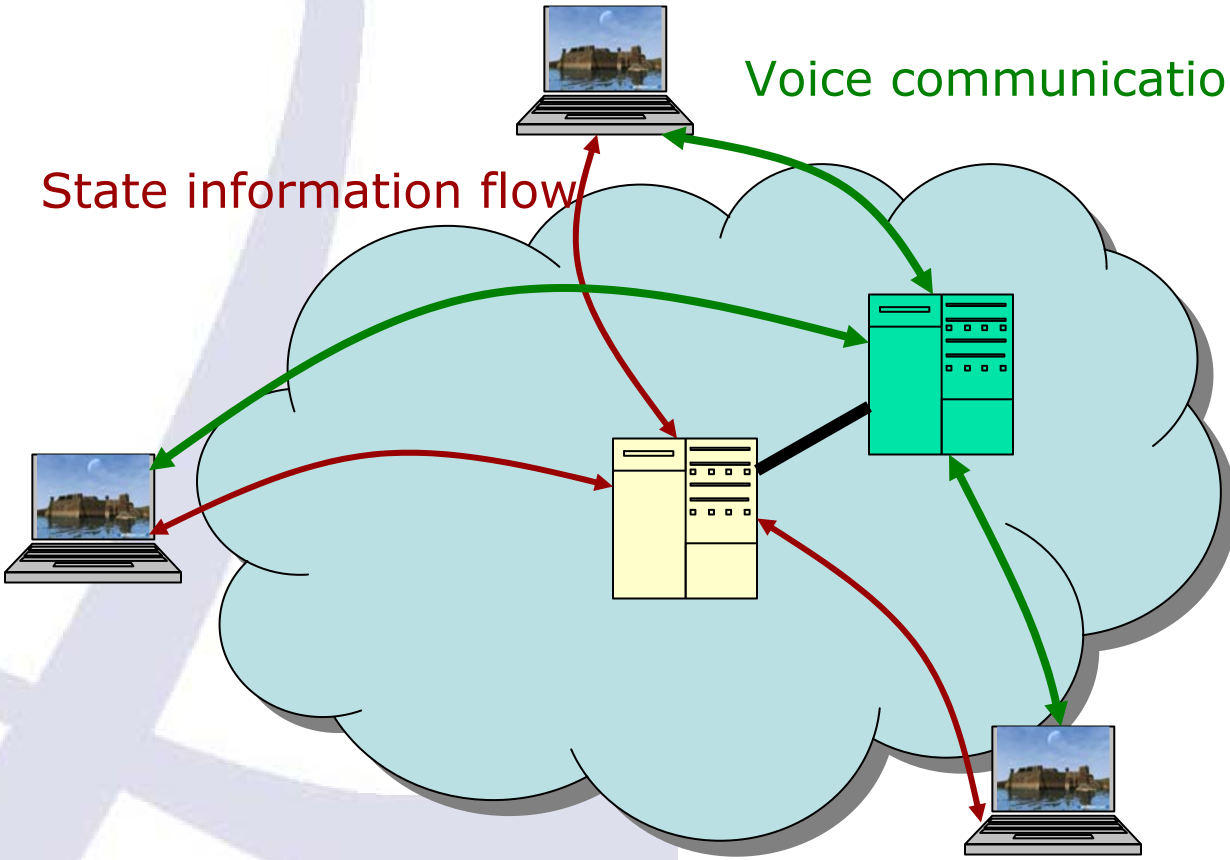

New Capabilities Required

- Latency driven distribution of processing in response to application dynamics
- Real time creation of content from a dynamic set of dispersed sources

Key Challenge (1)

Managing Latency

Response time

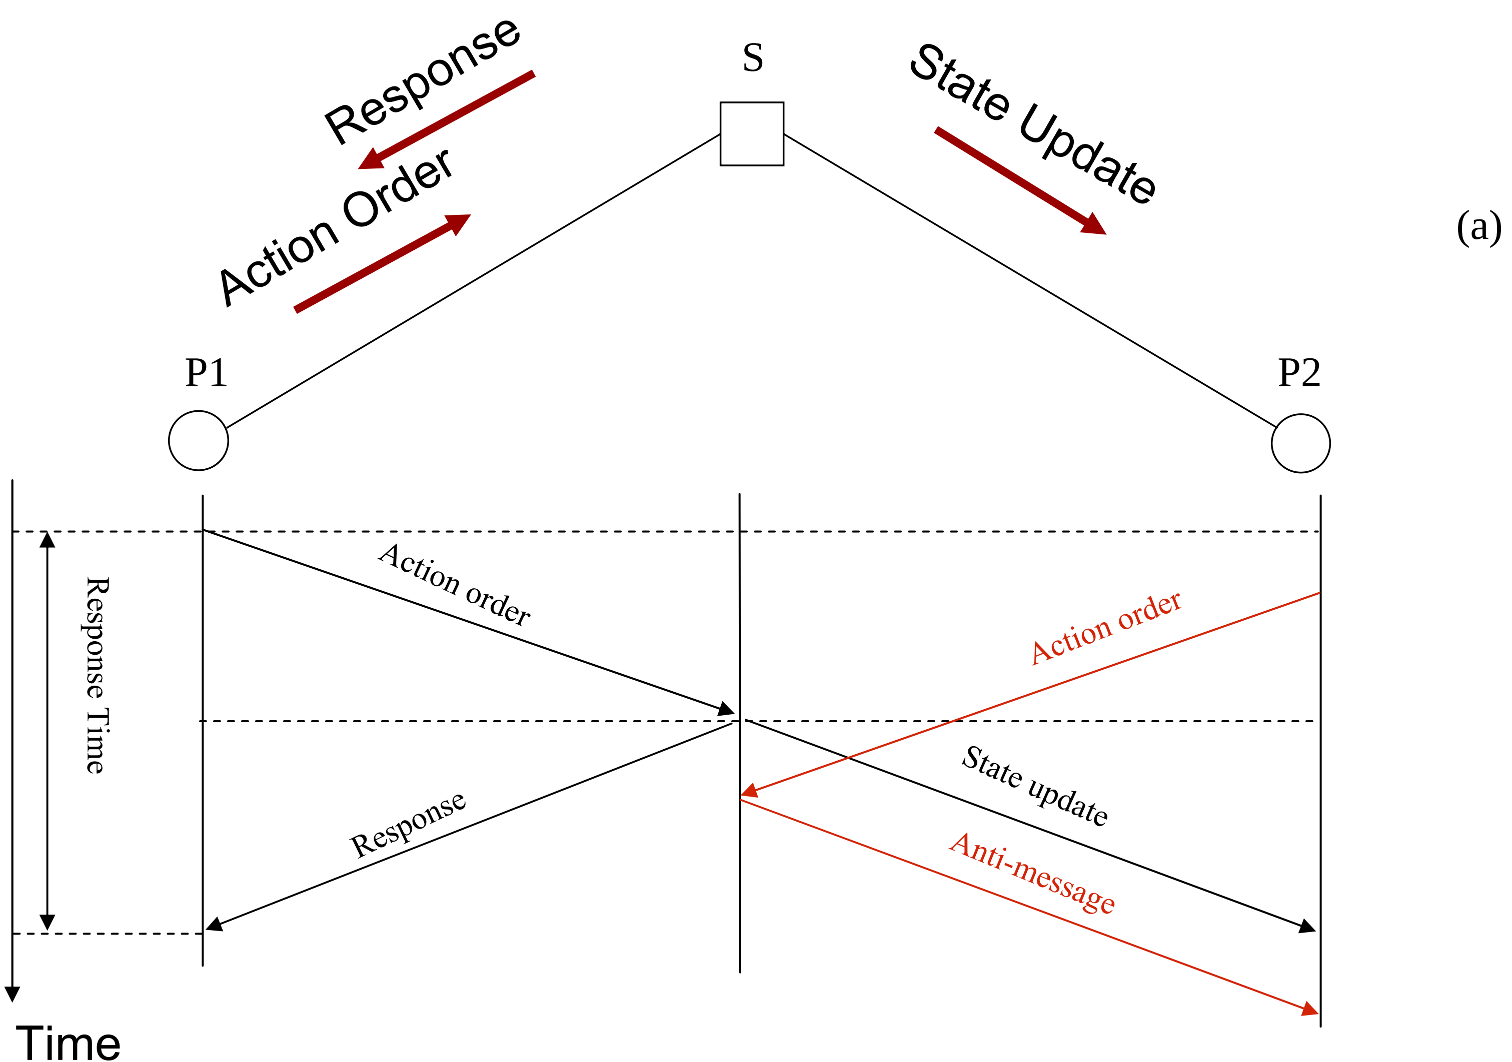

Response time after distribution

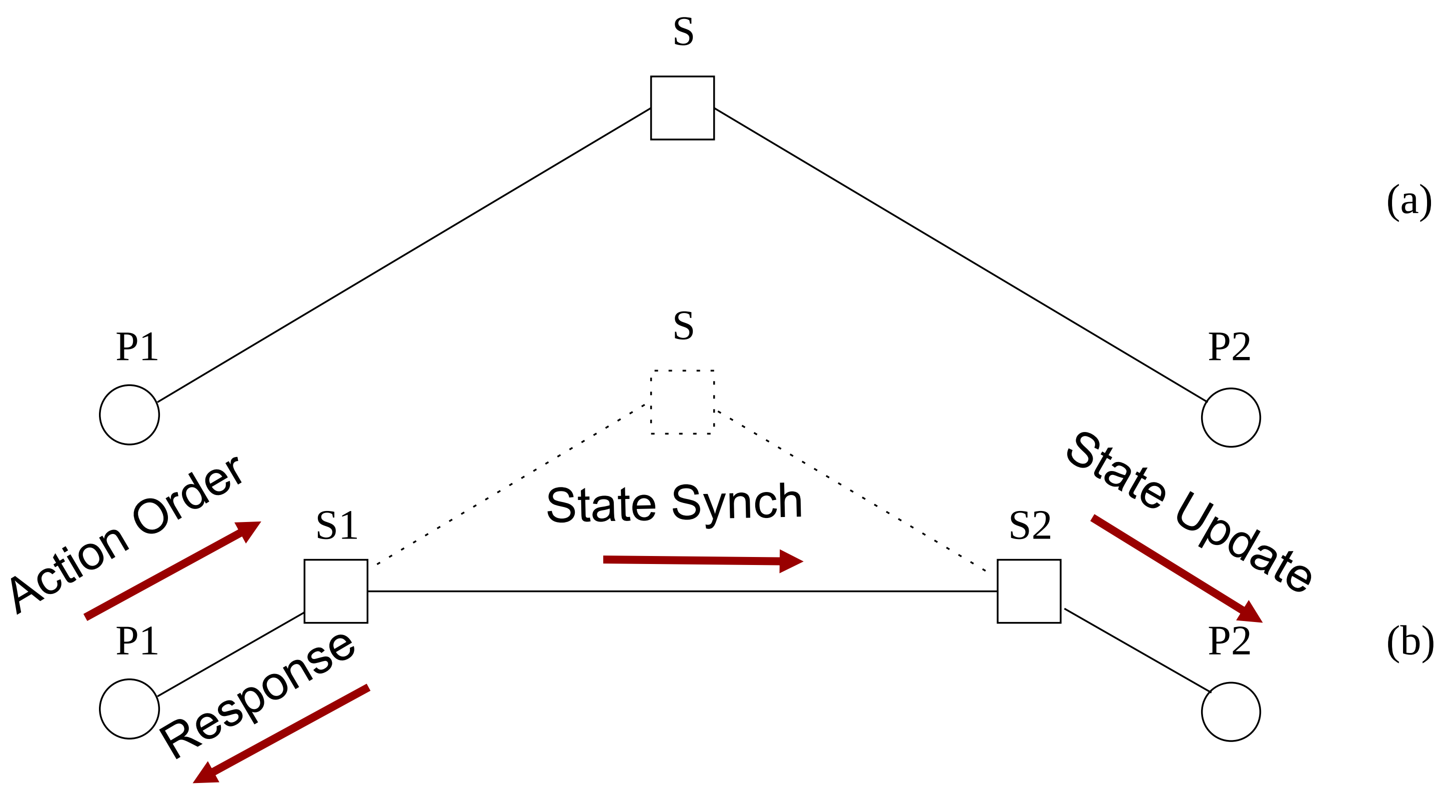

Paradoxes

- Healing
  - Rollback
- Avoidance
  - Server Local Lag

Local Lag

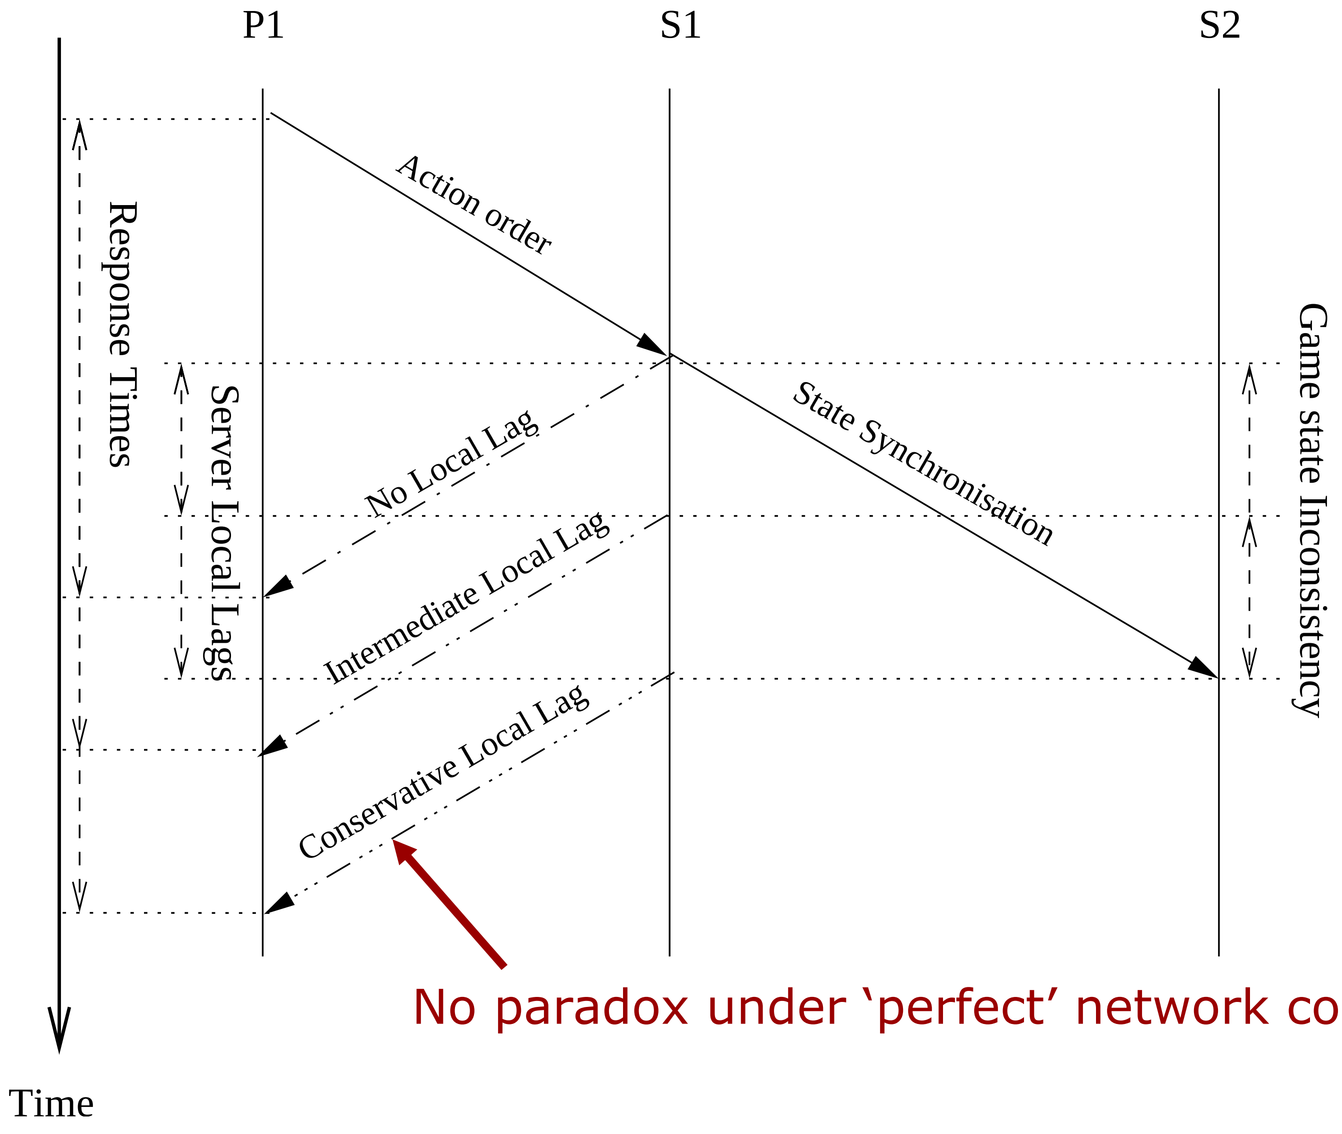

No paradox under 'perfect' network conditions

Improve Response Time

- Distributing the servers
- Tuning Local Lag
- Unbinding state parameters
- Control network delay between servers

Tuning Local Lag

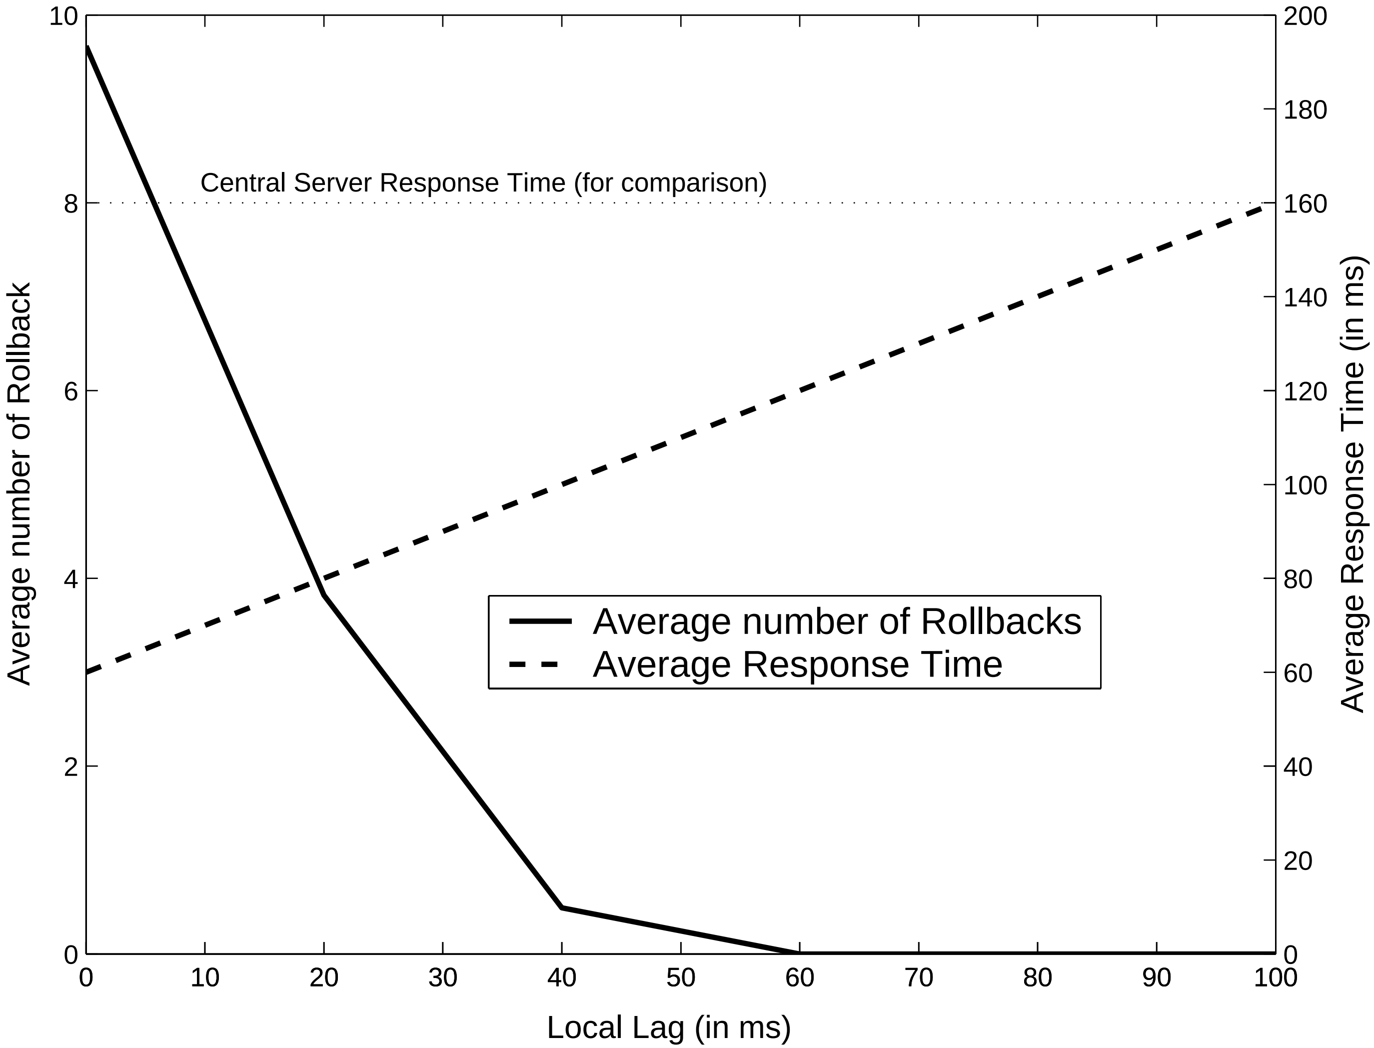

Rollback contours

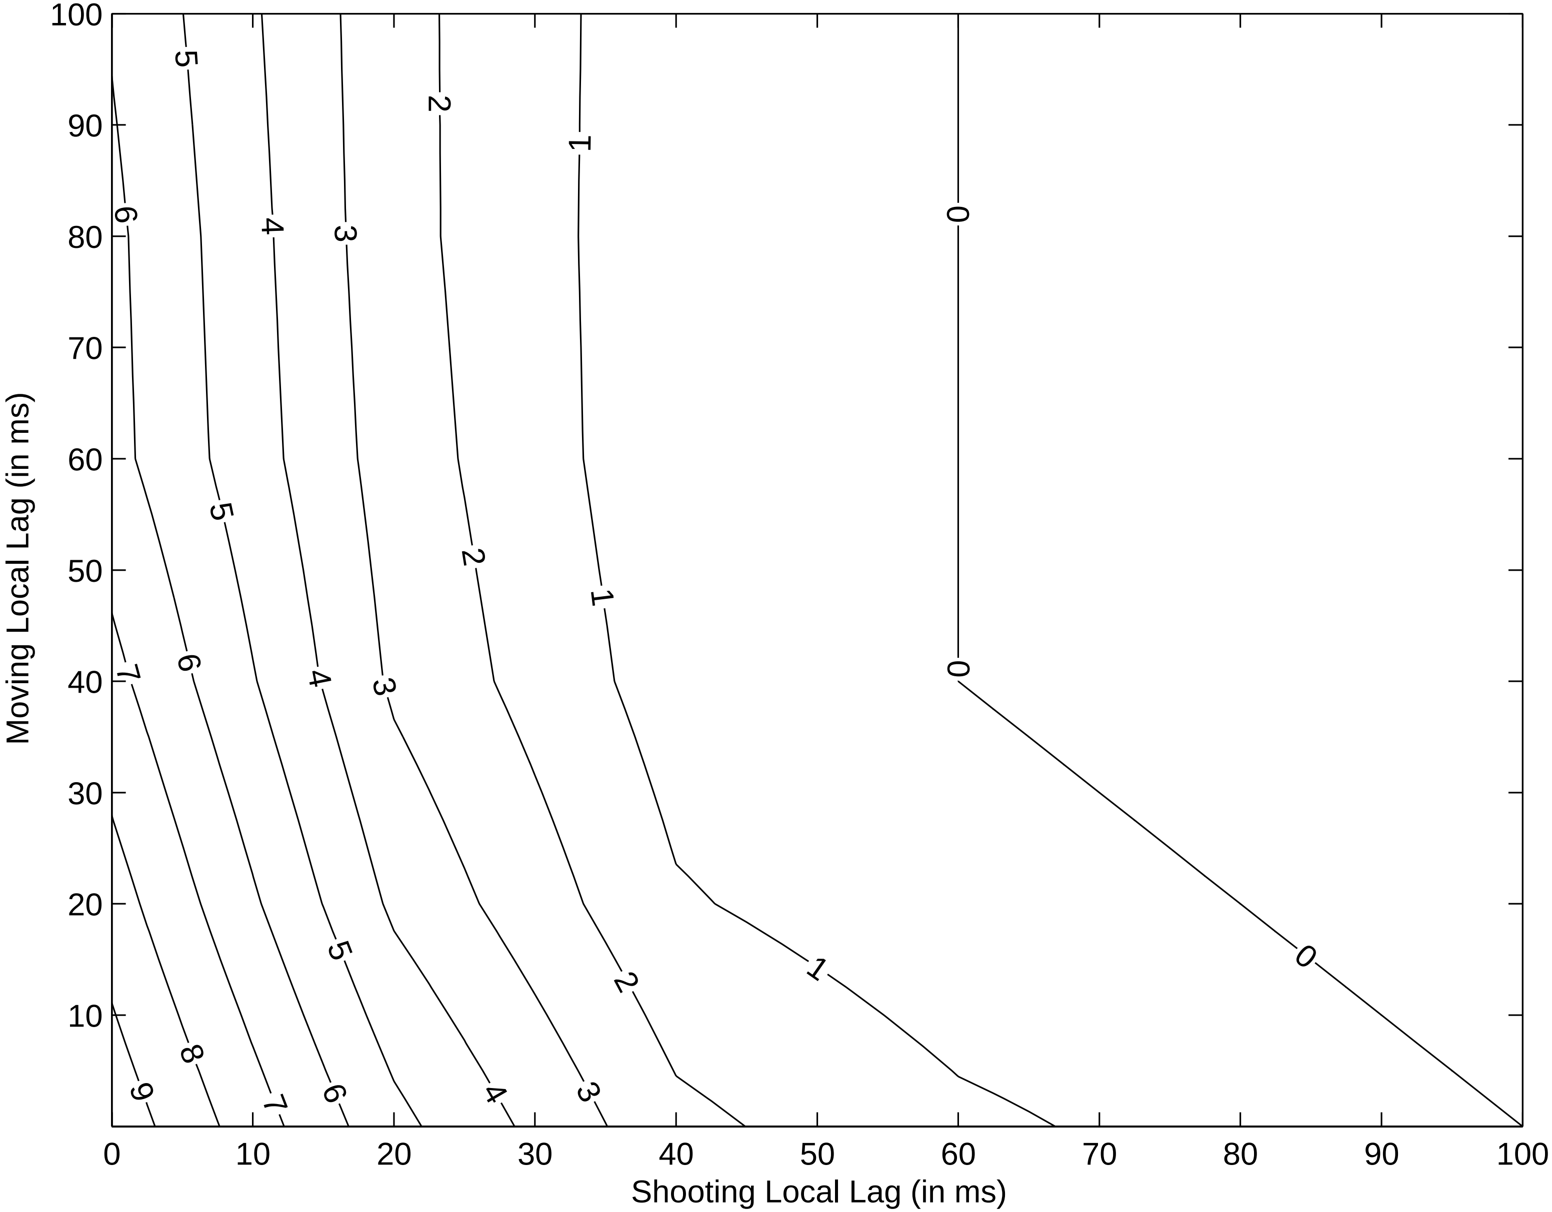

Response time contours

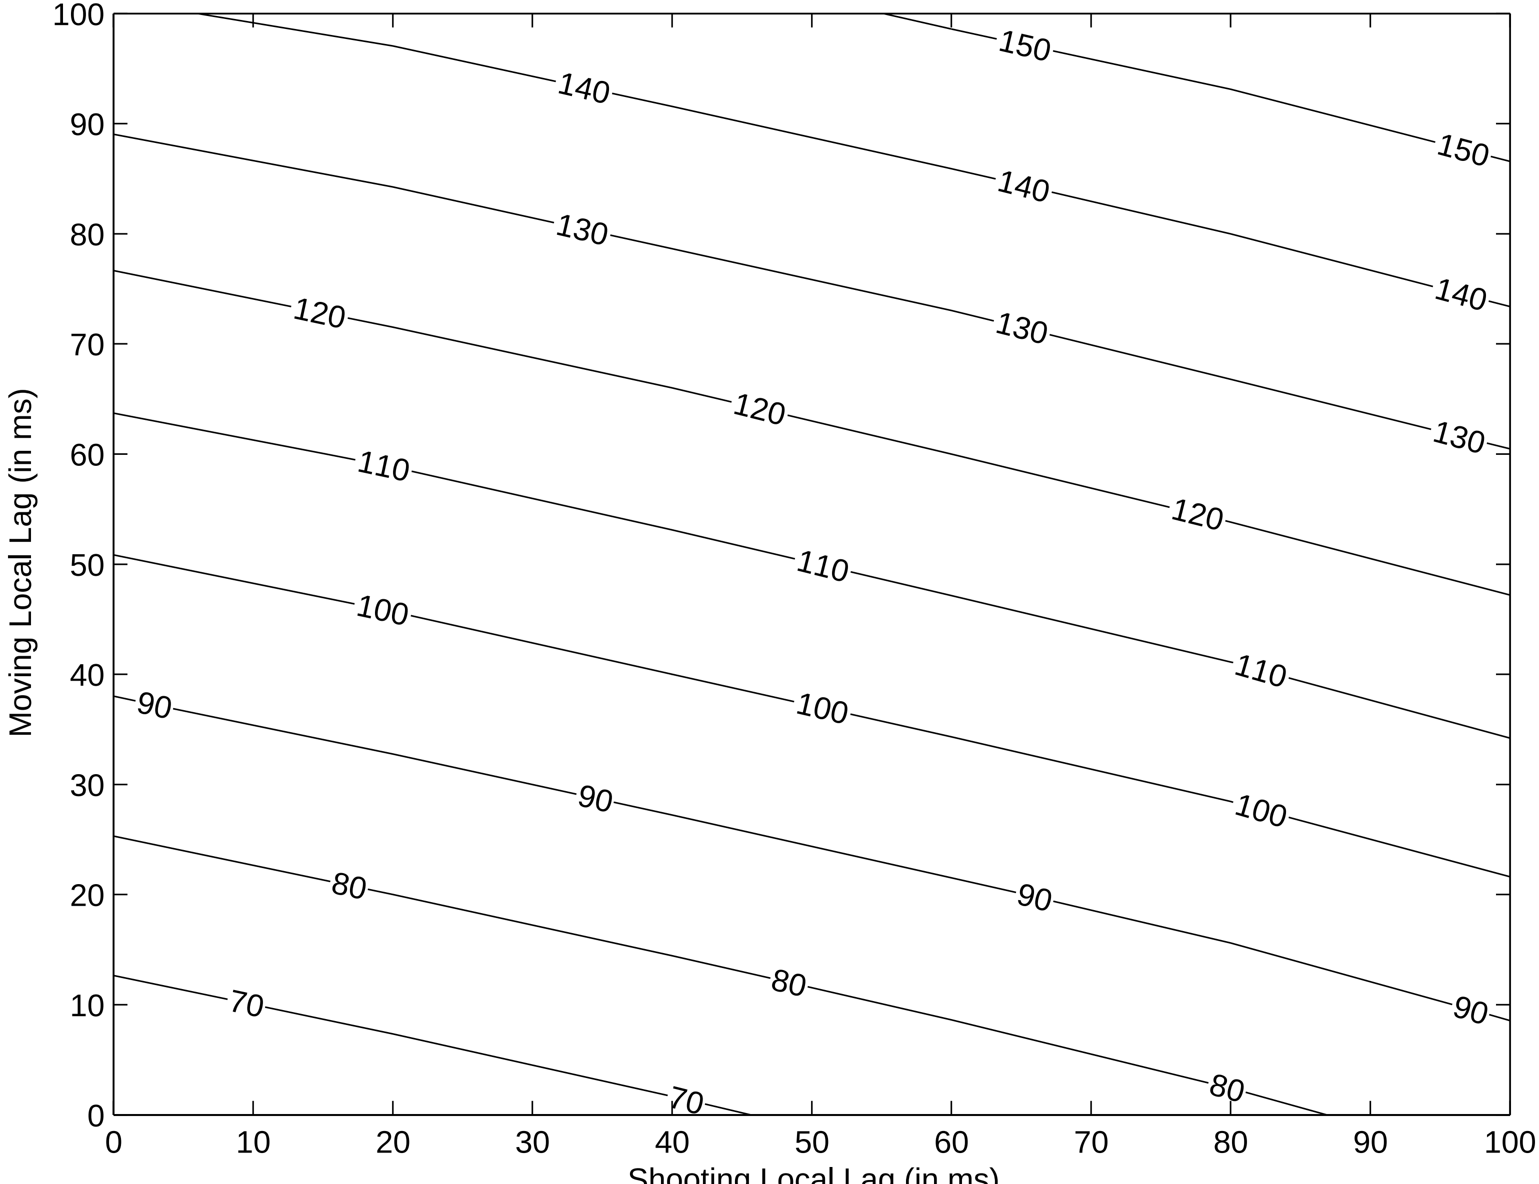

The Two Drivers for Distribution

- Resource Driven Distribution
  - Move components of an application to different processors when hitting the resource limit in one location
  - No performance gain due to distribution if resources available in one location

The Two Drivers for Distribution

- Latency Driven Distribution
  - Distribution will be needed even if no resource limitation exists
  - The primary performance parameter of interest is latency/responsiveness

Latency Driven Distribution

- It is important to
  - control where the processing is done, and
  - change that location dynamically

Key Challenge (2)

Real-time Creation of Content

Dense Immersive Communication Environment - DICE

- DICE creates a realistic voice communication for a virtual crowd
  - The voice of each member of crowd is heard in perfect harmony with his/her perceived spatial location
- Immersive and scalable to dense virtual spaces

Audio Server

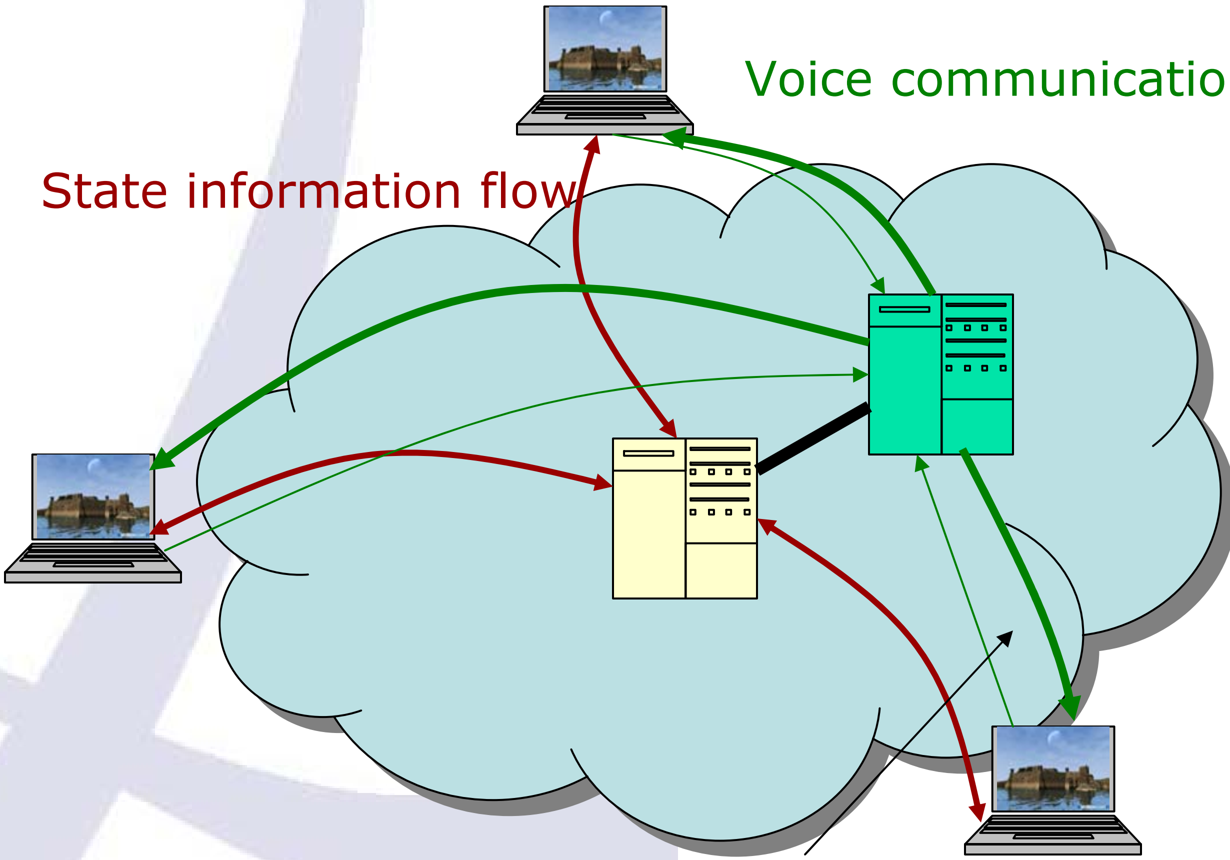

Improvement of distributed server relative to an optimal central server

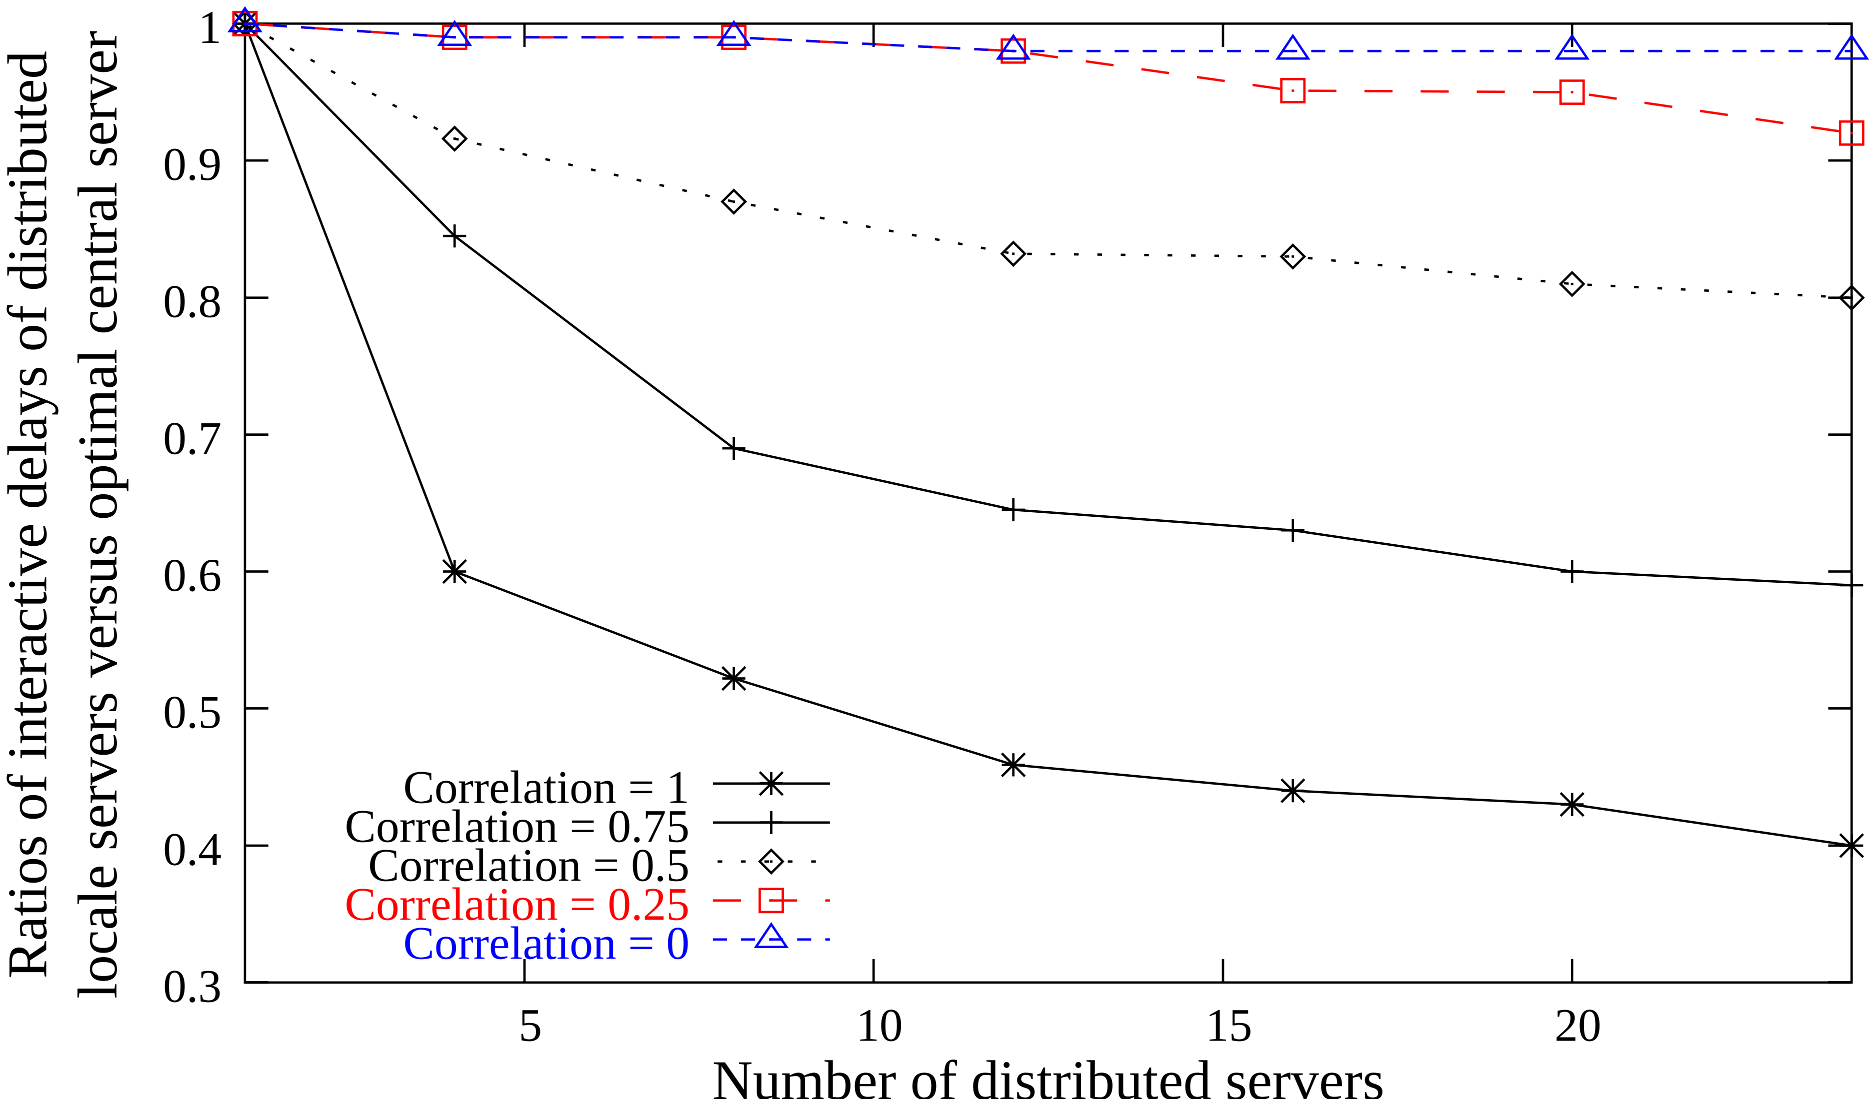

Infrastructure Design

LDD supporting infrastructure

- LDD only makes sense over large-scale infrastructure
  - Cannot ignore the network
- Applications need to control
  - Location of processing
  - Inter-server delay
  - Routing of flows
  - Packet level duplication, re-routing, multicast, …

LDD supporting infrastructure

- Need some ‘middleware’ to simplify application development
- Our ‘middleware’ is combination of software + hardware

LDD supporting infrastructure

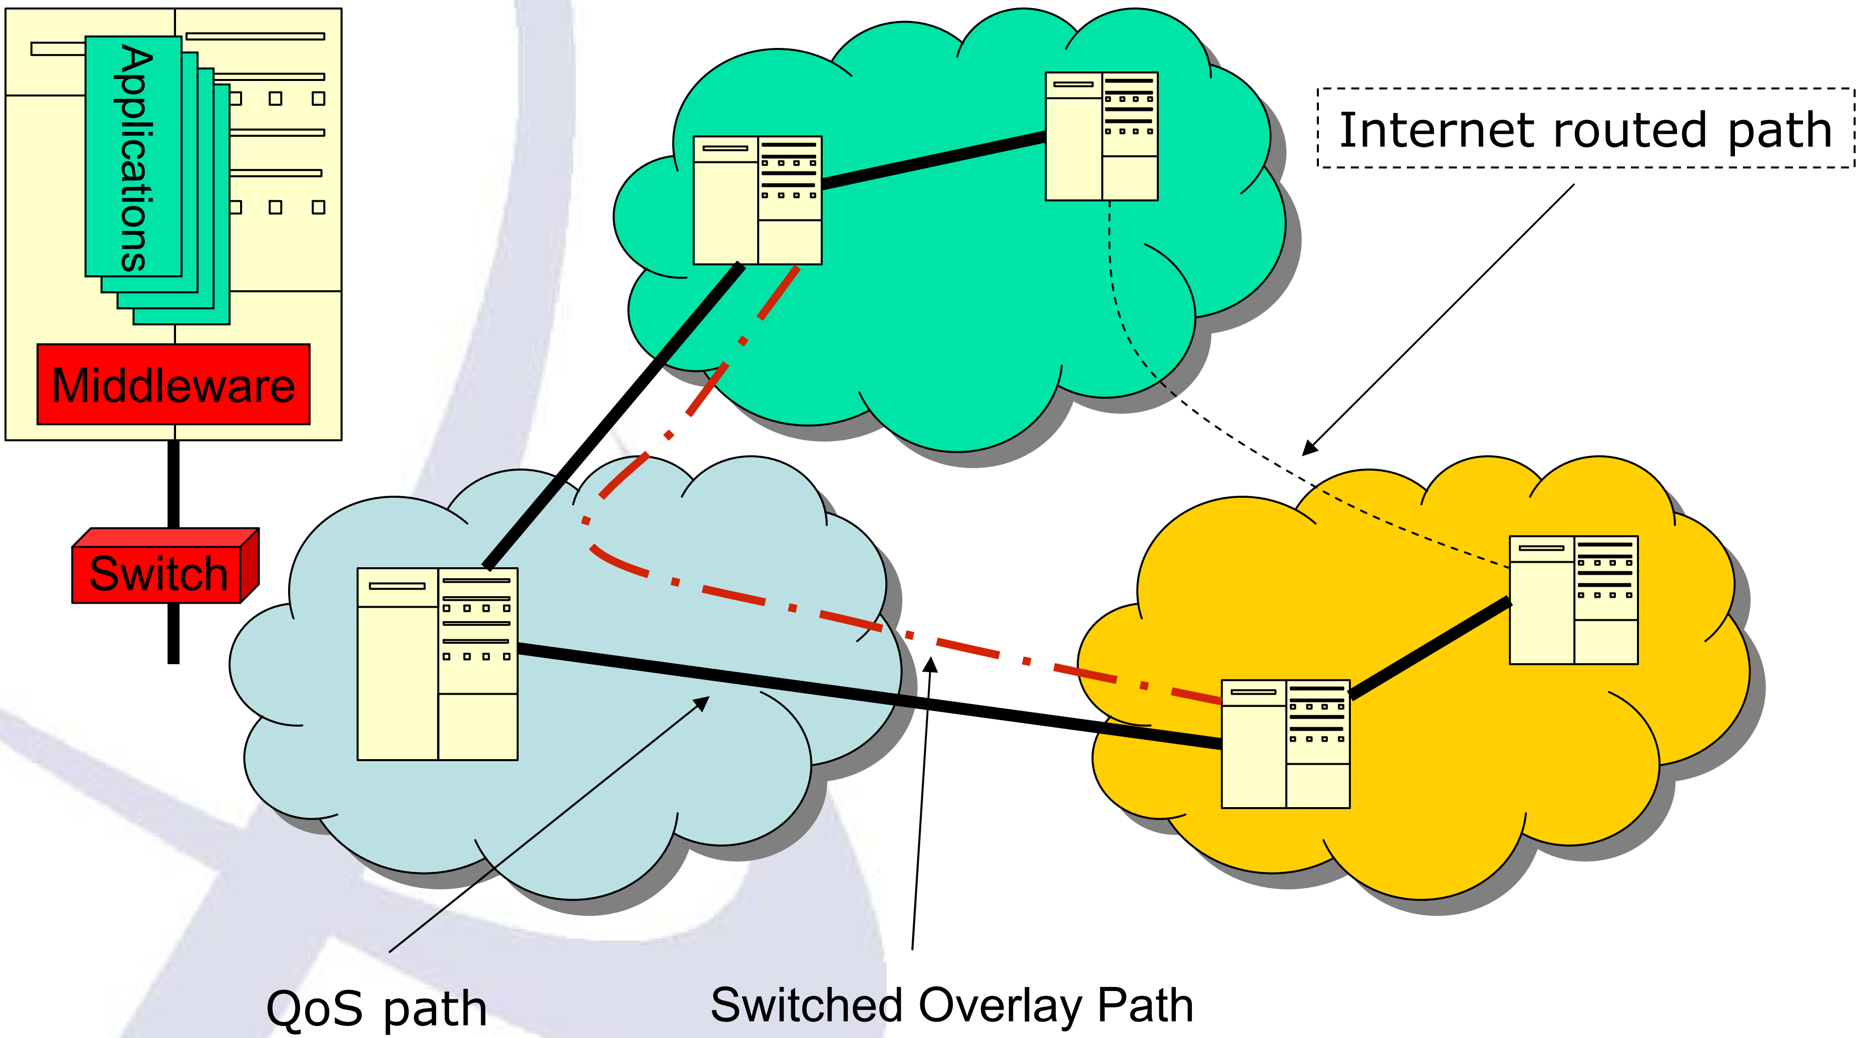

Switched Overlay Network (SWON)