Future Directions

1. Introduction
2. Fundamentals and design principles
3. Network architecture and topology
4. Network control and signalling
5. Network components
   - 5.1 links
   - 5.2 switches and routers
6. End systems
7. End-to-end protocols
8. Networked applications
9. Future directions

Future Directions
Looking toward the future

Future Directions
Know the Future

- In 1990
  - the Internet was for academics to email and FTP documents
  - the Web didn't exist (but the pieces did: hypertext, DVSM, xftp)
  - it was just becoming clear the the Internet *was* the GII
    - and any attempts to replace the IP waist are futile (e.g. ATM)
  - laptop computers were a novelty of only a few road warriors
    - and only dialup access available to home network

Know the Future $Ø_3$

The future hasn't happened yet, and is guaranteed to contain at least one completely unexpected discovery that changes everything.

Future Directions
Prepare for the Future

- We can't predict the future...

  ...but we can be *prepared*

Prepare for the Future Ø$_4$

Simply knowing the past does not prepare us to understand the future. We must constantly reevaluate tradeoffs in the face of emerging technology, and question the basic assumptions of the present.

Future Directions
Resource Tradeoffs Drive the Future

- Nonuniform advances in technology leads to relative shifts in the cost and availability of resources
  - P: processing
  - E: energy
  - M: memory
  - L: latency
  - B: bandwidth
- Understanding these shifts helps us prepare for the future

Resource Tradeoffs Change $Ø_4$

The relative cost of resources and the impact of constraints change over time, due to nonuniform advances in different aspects of technology.

Resource Tradeoffs
Shifts in Relative Costs

- First generation networks (through 1970s)
  - bandwidth scarce and expensive
  - users were careful with network interaction
- Third generation networks (1990s)
  - bandwidth became cheap relative to processing
  - bottlenecks in store-and-forward routers and end-systems
  - Web drove bandwidth demand and changed interaction
    - network looks like a big timesharing system
    - client/server transactions
  - peer-to-peer networking applications were *not* a surprise
    - although network providers engineered asymmetric networks

Resource Tradeoffs
Fundamental Changes in Relative Costs

- Fundamental changes in tradeoffs are harder to see
- What if...
  - OC-3072 to the desktop? ...everything else constant
  - 1THz computer on the desktop? ...everything else constant
  - 1PB disk on the desktop? ...everything else constant

Potential Future Direction
Network Infrastructure: Optical

- Optical networking
  - significant advances in technology have occurred
    - EDFAs enabled long haul optical links *overnight*
    - MEMS switching elements
    - free space photonic links
  - other seem elusive
    - wavelength converters
    - photonic header processing
  - others are almost certain to occur
    - high-power steerable free space laser links

Potential Future Direction
Network Infrastructure: Wireless

- Mobile wireless networking
  - significant demand is already driving...
  - but beware of deployment disasters: Iridium, WAP, 3G/4G
  - industry *not* leading toward ubiquitous personal networking
    - e.g. Bluetooth 3-bit address space debacle
    - 802.11 + Mobile IP not even close to future needs
    - jury out on IETF MANET, but future not promising
- Efficient wireless–optical boundary?
  - unreliable packet-based wireless access
  - reliable circuit-based optical core
  - wireless burst switching?
  - optical CDMA?

Potential Future Direction
Network Infrastructure: Services

- Internet architecture increasingly disorganised
  - network service provider and enterprise hacks
  - middleboxes: NATs, firewalls, etc
- User and application overlays
  - multicast, peer-to-peer file sharing
  - not translucent; no knobs and dials
- Networks difficult to use; impossible to manage
  - extensible network framework needed
  - *systematic* mechanisms needed for service deployment
  - network service providers should want moderate active nets
  - overall application-to-application performance should matter

Potential Future Direction
Application Scenarios

- Ubiquitous computing and smart spaces
  - tera- and peta-node networks
    - hundreds to millions of devices per user
    - quadrillions to quintillions of sensor/actuator nodes
- Teleimmersion and distributed virtual reality
  - what is the bandwidth of a distributed holodeck?
- Distributed computing and storage networks
  - from distributed information to distributed computing
- Interplanetary and intergalactic Internet
  - RTT to Pluto 10^{10} km is 20 hours
  - bandwidth-×-delay product at 1Tb/s is 36 Pb

Future Directions
Conclusion

9.1 Looking toward the future

9.2 Conclusion

High-Speed Networking is Mature

- Everything has *some* aspect of high-speed...
  ...nothing is *only* high speed

- Difficult to
  - maintain a distinct high-speed networking community
  - justify funding for high-speed networking research

The Disorganisation of the Net

- Once possible to understand networks as a whole
  - PSTN (before deregulation, local number portability)
    - deregulation not *controlled* – resulted in chaos
  - hierarchical NSFNET before Internet privatisation
    - privatisation not *controlled* – resulted in chaos
- The current network is a mess...
  - topology
    - multiple service providers arbitrarily peered
    - transparent overlays without ability to convey performance
  - middleboxes and ASP hacks
    - end-to-end semantics hopelessly gone
    - no concern for overall impact of individual hacks

Role for the High-Speed Community?

- Who is looking out for the performance as a whole?
  - end-to-end
  - application-to-application
- Can we have application-to-application high-speed?
  - high bandwidth: perhaps
    - technology driven (silicon and optics)
  - low latency: less likely
    - speed-of-light constraints
    - latency and inefficiency through hacks