High-Speed Networking:
A Systematic Approach to High-Bandwidth Low-Latency Communication

James P.G. Sterbenz
University of Kansas EECS/ITTC
Lancaster University Computing Dept.
jpgs@sterbenz.org

http://hsn-book.sterbenz.org

+1 508 944 3067

Abstract

This tutorial presents a comprehensive introduction to all aspects of high-speed networking, based on the book *High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*, James P.G. Sterbenz and Joseph D. Touch, John Wiley, 2001. The target audience includes computer scientists and engineers who may have expertise in a narrow aspect of high-speed networking (such as switch design), but want to gain a broader understanding of all aspects of high-speed networking and the impact that their designs have on overall network performance. This tutorial is not about any particular protocols and standards, but is rather a systemic and systematic approach to the principles that guide the research and design of high-speed networks, protocols, and applications.

The network is a complex system of systems, and high-speed networking does not result from the design of individual components or protocols in isolation. Thus, this tutorial presents a systemic approach to high-speed networks, where the goal is to provide high bandwidth and low latency to distributed applications, and to deal with the high bandwidth-×-delay product that results from high-speed networking over long distances. A set of fundamental axioms is presented (Know the past present and future, Application primacy, High-performance paths, Limiting constraints, and Systemic optimisation) followed by the major topics: Network architecture and topology; Network control and signalling; Communication links; Switches and routers; End systems; End-to-end protocols; Networked applications.

A set of design principles are defined and applied to each of the topics:

1. Selective optimisation; 2. Resource tradeoffs; 3. End-to-end arguments; 4. Protocol layering; 5. State management; 6. Control mechanism latency; 7. Distributed data; 8. Protocol data unit structure.

A set of design techniques (scaling time and space, masking the speed of light, specialised hardware implementation, parallelism and pipelining, data structure optimisation, cut-through and remapping) are introduced and applied as appropriate.

Revision History

1.0	IEEE Hot Interconnects,	Stanford,	23 Aug. 2002
1.1	IEEE Networks 2002,	Atlanta,	27 Aug. 2002
2.0	IEEE Hot Interconnects,	Stanford,	22 Aug. 2003
2.1	IEEE MILCOM,	Boston	15 Oct. 2003
2.2	Universität der Bundeswehr München		15 Mar. 2004
2.3	IEEE Hot Interconnects,	Stanford,	27 Aug. 2004
2.4	IEEE Hot Interconnects,	Stanford,	17 Aug. 2005
2.5	IEEE Hot Interconnects,	Stanford	25 Aug. 2006

Sources

This tutorial is based on...

- James P.G. Sterbenz and Joseph D. Touch, *High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication*, John Wiley, New York NY US, 2001
  - with contributions from
  - Julio Escobar
  - Rajesh Krishnan
  - Chunming Qiao
  - A. Lyman Chapin

Introduction

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

Scope

- All factors needed for high-speed networking
  - network components
  - protocols
  - network: a complex system of systems
  - end-to-end delivery of data to applications
  - applications that use and drive high-speed networks
- Lots of networking topics are *not* covered
- Ask questions throughout!

What is High Speed?
Bandwidth and Latency

- Delay
  - D end-to-end
  - d per hop
- Bandwidth
  - B aggregate
  - b per flow
- Bandwidth-×-delay product
  - number of bits in flight on a high-speed path
  - b [bits/sec] × d [sec] = [bits]

Tutorial Structure

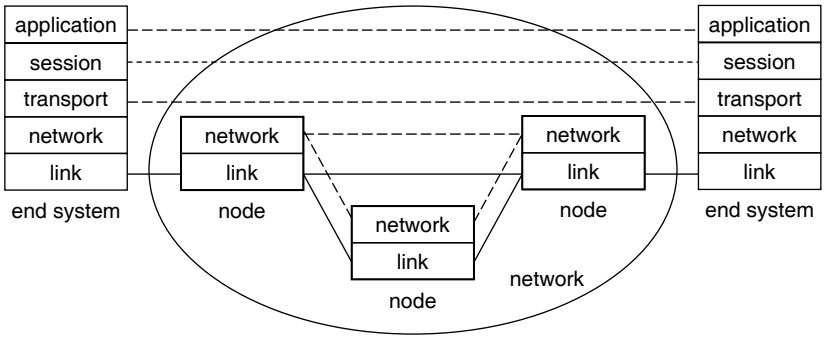

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Network Architecture and Topology

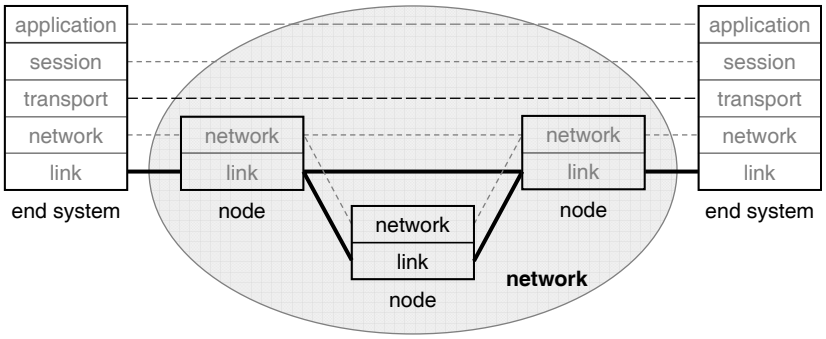

3.1. Topology and geography
3.2. Scale
3.3. Resource Tradeoffs

Network Control and Signalling

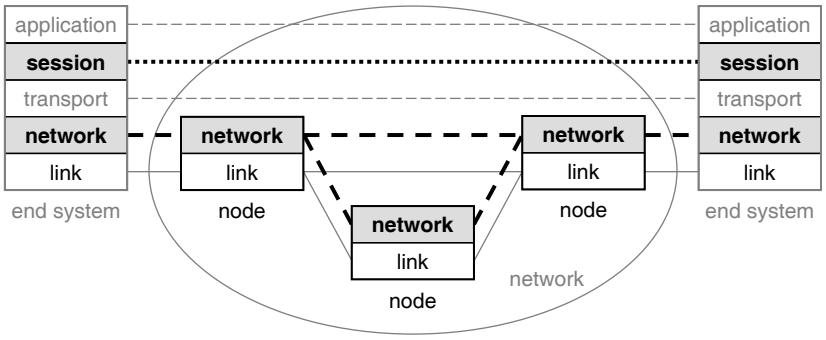

4.1. Signalling and control
4.2. Traffic management
4.3. Path routing dynamics
4.4. Monitoring & management

Network Components: Links

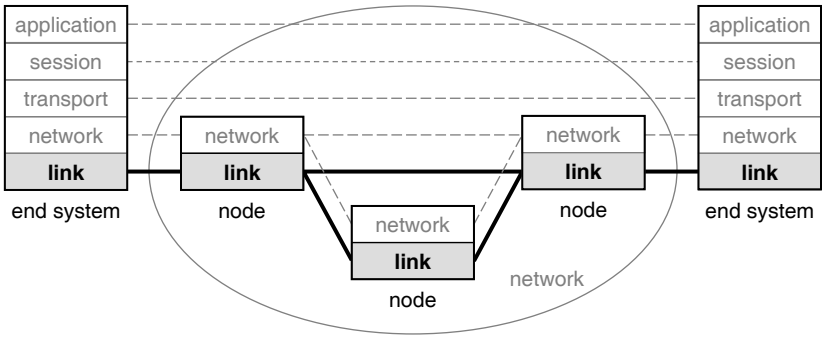

5.1.1. Physical transmission
5.1.2. Link technologies
5.1.3. Link layer components
5.1.4. Support for higher layers

Switches and Routers

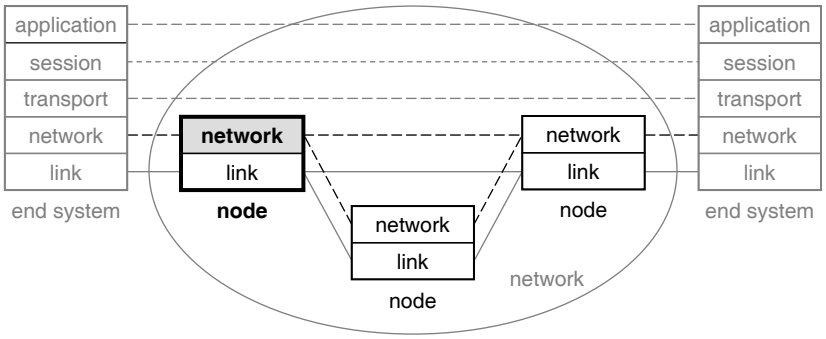

5.2. Switches and routers
5.3. Fast packet switches
5.4. Switch fabric architecture
5.5. Fast datagram switches
5.6. Higher layer and active processing

End Systems

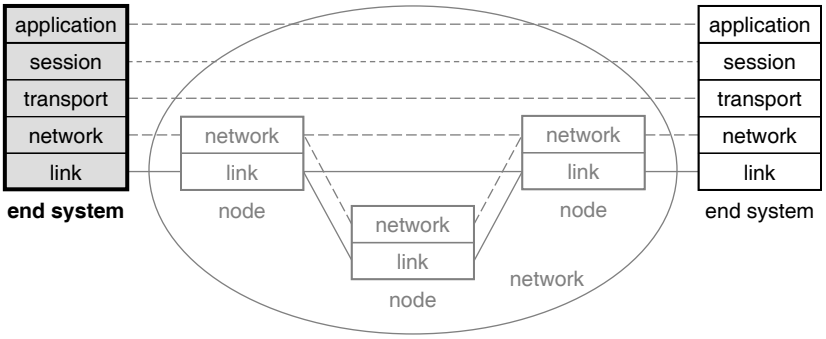

6.1. End system components
6.2. Protocols and OS software
6.3. End system organisation
6.4. Host–network interface

End-to-End Protocols

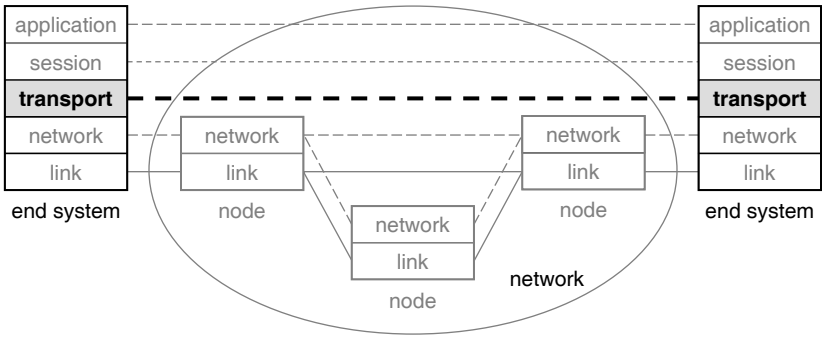

7.1. Functions and mechanisms
7.2. State management
7.3. Framing and multiplexing
7.4. Error control
7.5. Flow & congestion control
7.6. Security & info assurance

Networked Applications

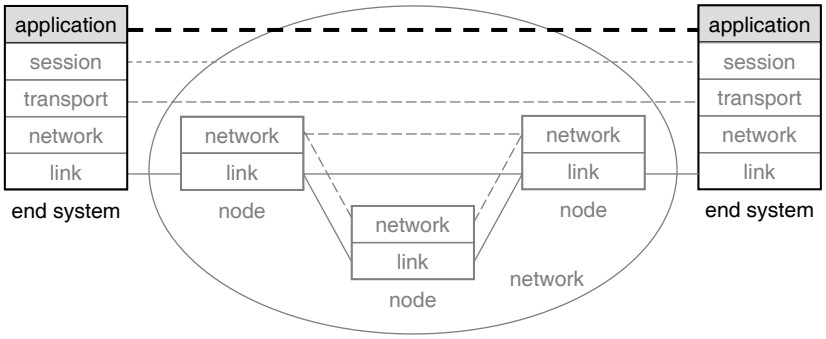

8.1. Application characteristics
8.2. Application categories
8.3. Application adaptation
8.4. Network interaction