CSC7601

High Speed Networks

Lonnie Leger
Louisiana State University
LONI Director of Networking
lonnie@lsu.edu
225-578-8391

http://www.loni.org

Abstract

The network is a complex infrastructure of systems, and high-speed networking does not result from the efforts of individual components in isolation from the others. The goal is to provide high bandwidth and low latency to grid applications in dealing with the high bandwidth-x-delay product that results from high-speed networking over long distances.

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

1. Selective optimization
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

Scope

- All factors needed for high-speed networking
  - network components
  - protocols
  - Complex system of systems
  - end-to-end delivery
  - applications
- Lots of networking topics are not covered
- Ask questions throughout!

Sources

- http://www.davesite.com/webstation/net-history.shtml
- http://www.zakon.org/robert/internet/timeline/
- www.nlr.net
- http://currents.ucsc.edu/06-07/11-27/brief-varma.asp
- http://www.sterbenz.org/jpgs/tutorials/hsn/
- http://williamstallings.com/HsNet2e.html

Table of Contents

Define

- Transmission capacity
- Protocol efficiency
- Cost
- Latency
- Application
- Distance
- Host

What is High Speed?

- Delay

  D end-to-end

  d per hop
- Bandwidth

  B aggregate

  b per flow
- Bandwidth-x-Delay product
  - number of bits in flight on a high-speed path

  b [bits/sec] x d [sec] = [bits]

Complex Infrastructure

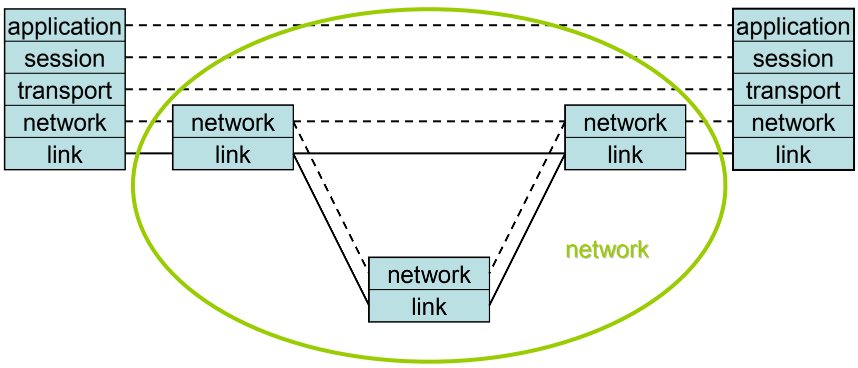

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Network Architecture

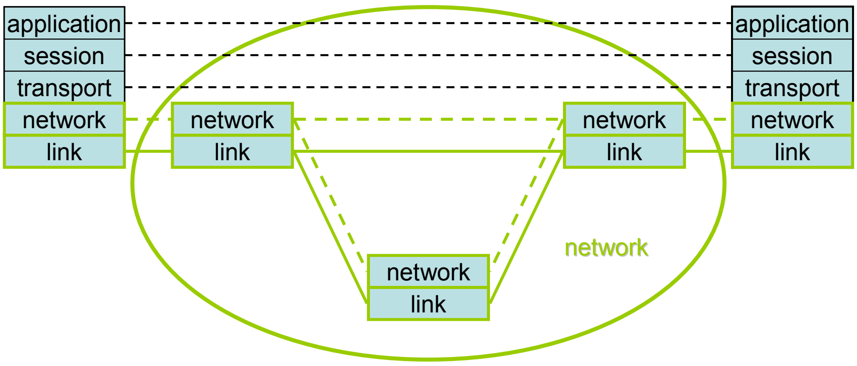

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Network Control and Signalling

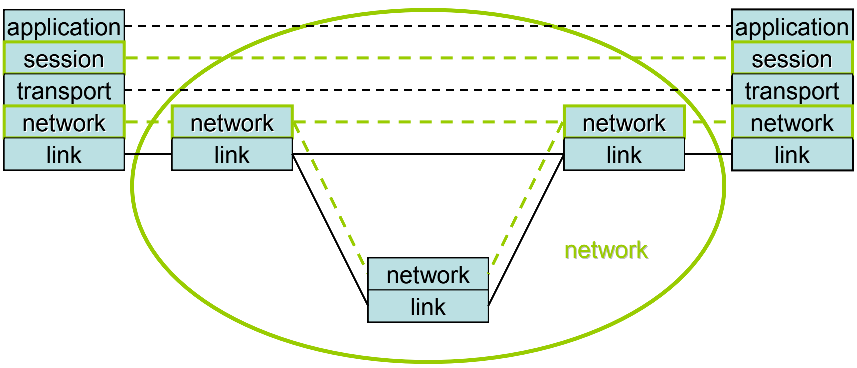

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Links

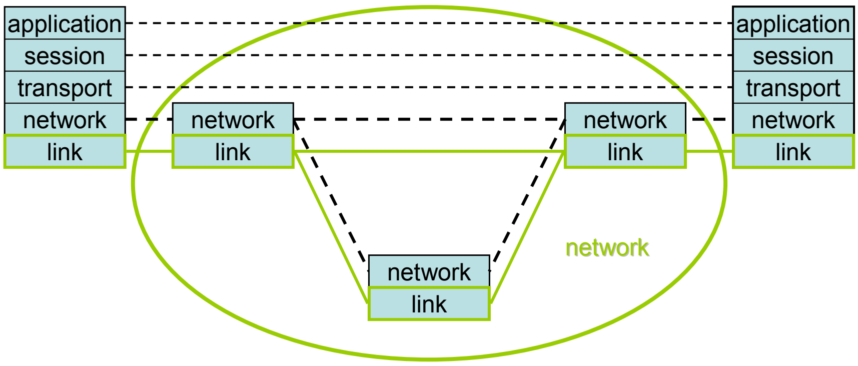

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

26 September 2007 CSC7601 11

Routers and Switches

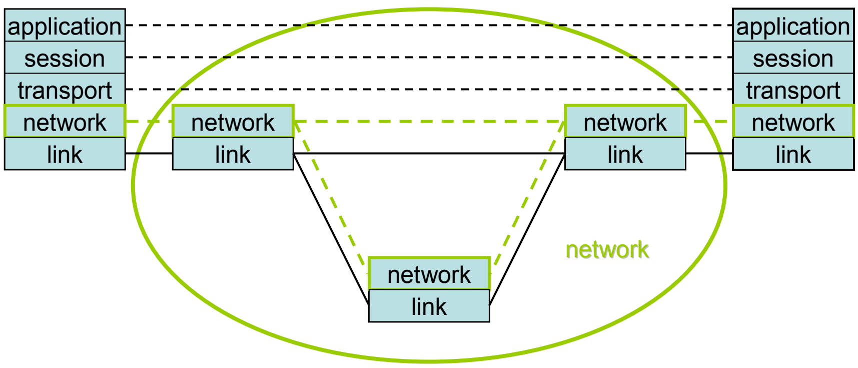

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

26 September 2007 CSC7601 12

End Systems

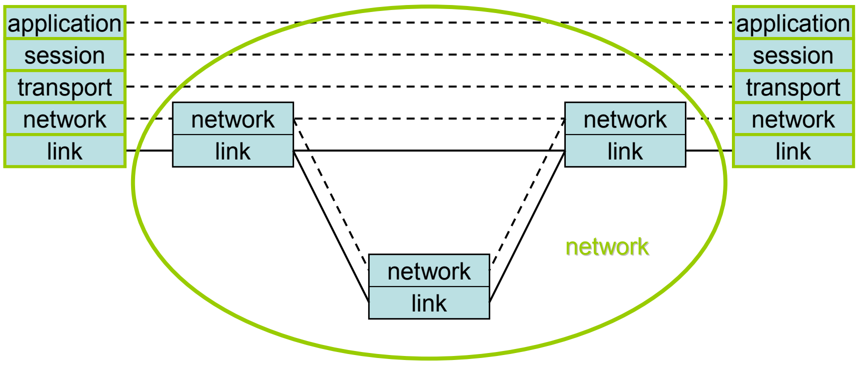

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

End-to-End Protocols

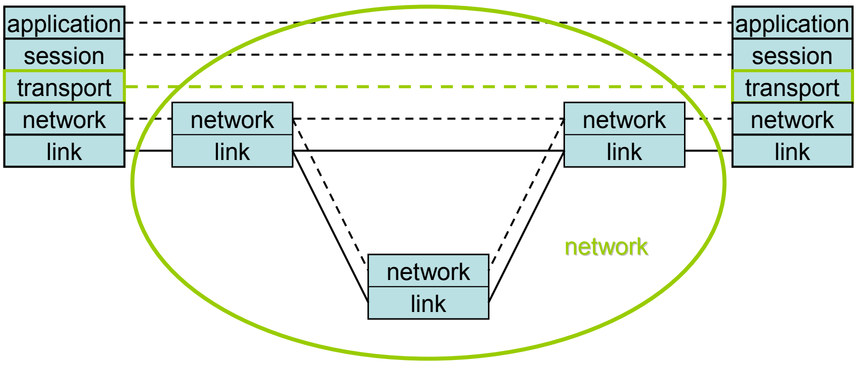

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Networked Applications

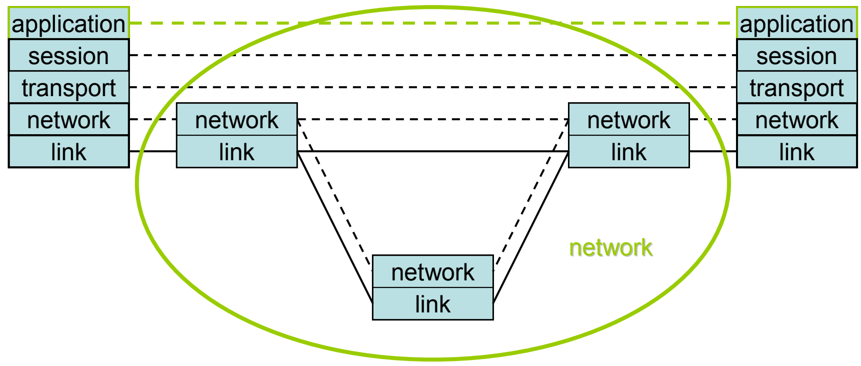

- Bottom up
  - network components
  - applications
- Inside out
  - network components
  - end systems

Table of Contents

Fundamental Reasoning

Know the Past, Present and Future

- It is necessary to know the past
  - to know what is really new
  - to not waste time reinventing the past
  - to not repeat past mistakes
- It is necessary to have a broad current understanding
  - to understand how to reapply past ideas
  - to know what new needs to be done
- It is necessary to be aware that the future will change everything

How we got Here

- 1968 – 50Kbps on ARPA with 4 hosts
- 1981 – 56Kbps on CSNET with 213 hosts
- 1985 – 1.544Mbps on NSFNET with 1961 hosts
- 1992 – 45Mbps on NSFNET with 1M+ hosts
- 1994 – 155Mbps on NSFNET with 3M+ hosts
- 1994 – NSFNET traffic passes 10 trillion bytes/month
- 1996 – 622Mbps on Internet
- 1999 – 2.5Gbps on vBNS
- 2004 – 10Gbps on Internet2
- 2004 – 10Gbps on NLR
- 2005 – 10Gbps on LONI
- 2006 – 100Gbps demonstration

Fundamental Reasoning

2.1 Know the past, present and future

2.2 Drivers and constraints

  2.2.1 Application primacy

2.3 Design principles and tradeoffs

2.4 Design techniques

Application Primacy

"The sole and entire point of building a high-performance network infrastructure is to support the grid applications that need it."

- Effect
  1. Field of dreams vs. killer app dilemma
  2. Inter-application delay
  3. Network bandwidth and latency
  4. Network importance is system design

Application Primacy

- Field of dreams vs. killer app dilemma
  - "The emergence of the next "killer application" is difficult without sufficient infrastructure."
  - Applications need infrastructure on which to build
  - Infrastructure deployment needs to justify expense
  - Difficult to resolve without government funding

Application Primacy

- Field of dreams vs. killer app dilemma
- Inter-application delay
  - "The Performance metric of primary interest to grid applications is the total delay in transferring data."
  - Users and applications care about delay
    - Not bandwidth! (directly)
    - Users: inter-application delay
    - Applications: end-to-end network and end system delay
    - No difference between remote and local

Application Primacy

- Field of dreams vs. killer app dilemma
- Inter-application delay
- Network importance is system design
  - "Bandwidth and latency are the primary performance metrics important to inter-application delay."
  - Inter-application delay *D* consists of two components
    - path latency d_{path}
    - transmission delay $d_{\text{transmission}} = b$ [bit] / r [bit/sec]
    - $D = d_{\text{path}} + d_{\text{transmission}}$
  - Low latency directly needed
  - High bandwidth needed based on object size

Application Primacy

- Field of dreams vs. killer app dilemma
- Inter-application delay
- Network bandwidth and latency
- Network importance is system design
  - "Communication is the defining characteristic of networked applications, and thus support for communication must be an integral part of systems support grid applications."
    - Communications performance should drive design
    - Should not be driven by consumer PC architecture

Fundamental Reasoning

2.1 Know the past, present and future

2.2 Drivers and constraints

- 2.2.1 Application primacy
- 2.2.2 High-performance paths goals

2.3 Design principles and tradeoffs

2.4 Design techniques

High-Performance Paths Goals

"The network and end systems must provide a low-latency, high-bandwidth path between applications to support low inter-application delay."

- Effect
  1. Path establishment – signalling, routing, and control
  2. Path protection – resource reservation or overprovisioning
  3. Store-and-forward and copy avoidance
  4. Blocking avoidance
  5. Contention avoidance
  6. Path information assurance tradeoff (against performance)

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
  - "Signalling and routing mechanisms must exist to discover, establish, and forward data along the high performance paths."
  - High performance paths need to be well-known
    - Signalling is needed to establish paths
    - Routing algorithms are need to determine path route
    - Forwarding mechanisms move data along path
  - Paths may need to be reconfigured
    - Routing around faults, congestion, opponents
    - In response to path dynamics: research

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
- Path protection – resource reservation or overprovisioning
  - “In a resource constrained environment, mechanisms must exist to provide high-performance path and protection during network congestion.”
  - High performance paths need to be protected
    - By overprovisioning
    - In a resource constrained environment
      - Resource reservation: network and end system
      - Congestion avoidance and control

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
- Path protection – resource reservation or overprovisioning
- Store-and-forward and copy
  - “Avoidance”
  - Per hop per byte delays are significant
  - Embrace cut-through or zero copy mechanisms

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
- Path protection – resource reservation or overprovisioning
- Store-and-forward and copy
- Blocking
  - "Avoidance"
  - Blocking delays packets
  - Embrace non-blocking routers and switches, alternate path routing, and steady-state queue length

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
- Path protection – resource reservation or overprovisioning
- Store-and-forward and copy
- Blocking
- Contention
  - "Avoidance"
  - MAC delays for shared medium
  - Embrace spatial reuse like parallel waves

High-Performance Paths Goals

- Path establishment – signalling, routing, and control
- Path protection – resource reservation or overprovisioning
- Store-and-forward and copy
- Blocking
- Contention
- Path information assurance tradeoff (against performance)
  - "Paths have application-driven reliability and security requirements, which may have to be traded against performance."
  - Control processing delays: authentication and keying
  - Data path delays: encryption/decryption and security headers

Fundamental Reasoning

Managing Constraints

"Real-word constraints make it difficult to provide high-performance path to applications."

- Effect
  1. Speed of light
  2. Channel capacity
  3. Attenuation and transmission power
  4. Routing and switching speed
  5. Cost and feasibility
  6. Heterogeneity
  7. Policy and administration
  8. Backward compatibility inhibits change
  9. Standards both facilitate and impede

Managing Constraints

- Speed of Light
  – "There is not direct optimization technique."
  – Propagation velocity are relatively fixed through different medium
  – Dictates fundamental limit on latency over a distance
  – Possible techniques to help meditate include: caching and prediction

Managing Constraints

- Speed of Light
- Channel Capacity
  - “The capacity of communication channels is limited by physics of the medium.”
  - Bandwidth of a medium
    - Dictates fundamental limit on data rate
    - Possible techniques to help meditate include: multiplexing and spatial reuse

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
  - “Attenuation of signals limits their propagation distance for a given transmission power.”
  - Attenuation limits the range of transmission
    - Guided: logarithmic behavior
    - Wireless: square law $1/r^2 - 1/r^4$
  - Transmission energy needed to compensate
    - Limited by transmitter power
    - Constrained by channel design parameters

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
  - "There are limits on switching frequency of network components."
  - Switching rate of electronic and photonic components
    - Dictates limit on data rate
    - Moore's law keeps reducing the constraint
    - Transistor complexity is proportional to data rate

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
- Cost and Feasibility
  - Competing solutions have different costs
  - Cost of deployment is significant determinant

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
- Cost and Feasibility
- Heterogeneity
  - Networks server to interconnect heterogeneous
  - Significant overhead needed to support heterogeneity

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
- Cost and Feasibility
- Heterogeneity
- Policy and Administration
  - "May constrain the paths through with applications can communicate and dictate how application functionality is distributed."
  - Economics and business models
  - Intellectual property and legal issues
  - Government regulation
  - Social dynamics

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
- Cost and Feasibility
- Heterogeneity
- Policy and Administration
- Backward Compatibility
  - "The difficulty of completely replacing widely deployed network protocols means that improvements must be backward compatible and incremental."

Managing Constraints

- Speed of Light
- Channel Capacity
- Attenuation and Transmission Power
- Routing and Switching Speed
- Cost and Feasibility
- Heterogeneity
- Policy and Administration
- Backward Compatibility
- Standards
  - "Standards are critical to facilitate interoperability but used too early can impede progress and too late become useless."

Fundamental Reasoning

Critical Path

"Optimize for the critical path, in both control and data plane."

- Optimize critical path on which high speed depends
  - Data transfer operations
  - Transfer control operations that assist in data transfer
    - Occur frequently
    - Cause serial delay to subsequent data transfer

Fundamental Reasoning

Resource Tradeoffs

"The relative composition of these resources must be balanced to optimize cost and performance."

- Effect
  1. Resource tradeoffs change the balance
  2. Optimal resource utilization vs. over engineering tradeoffs
  3. Support for multicast

Resource Tradeoffs

- Resource Tradeoffs Change
  - Relative cost of resources change over time
    - Non-uniform advances in technology
    - Changing application constraints
  - Design and engineering for the future
    - Track trends and shifts in resource tradeoffs
    - “Be Prepared”

Resource Tradeoffs

- Resource Tradeoffs Change
- Optimality vs. Over Engineering
  - Optimal bandwidth reservation and utilization costs
  - Optimize entire system cost

Resource Tradeoffs

- Resource Tradeoffs Change
- Optimality vs. Over Engineering
- Multicast
  - Needed by network control protocols
  - Important mechanism for conserving bandwidth

Fundamental Reasoning

2.1 Know the past, present and future

2.2 Drivers and constraints

2.3 Design principles and tradeoff

- 2.3.1 Critical path
- 2.3.2 Resource tradeoffs
- 2.3.3 End-to-end vs. hop-by-hop
- 2.3.4 Protocol layering
- 2.3.5 State and Hierarchy
- 2.3.6 Control mechanisms
- 2.3.7 Distribution of application date
- 2.3.8 Protocol data units

2.4 Design techniques

End-to-End vs. Hop-by-Hop

"Function required by communications applications can be correctly and completely implemented only with the knowledge and help of the applications themselves. Providing these functions as features within the network itself is not possible."

- Effect
  1. End-to-end argument
  2. Hop-by-hop performance enhancement

End-to-End vs. Hop-by-Hop

- End-to-End Argument
  - Hop-by-Hop functions do not compose end-to-end
    - Between HBH boundaries, function f is defeated
      - i.e. error control: errors may occur within components
      - i.e. encryption: cleartext within components may be snooped
      - These functions must be done E2E
        - » Doing them HBH is redundant, and may lower performance
  - Effect
    - Hop-by-hop performance enhancement
    - Endpoint recursion

End-to-End vs. Hop-by-Hop

- Hop-by-Hop performance enhancement
  - "It may be beneficial to duplicate an end-to-end function hop-by-hop, if doing so results in an overall (end-to-end) improvement in performance."
  - HBH functions should be duplicated if overall E2E benefit

End-to-End vs. Hop-by-Hop

- Endpoint Recursion
  - "What is hop-by-hop in one context may be end-to-end in another."
  - End-to-end arguments
    - Apply edge-to-edge
    - Subnet-to-subnet
    - Link-to-link
  - Overall E2E concerns must not be forgotten

Fundamental Reasoning

2.1 Know the past, present and future

2.2 Drivers and constraints

2.3 Design principles and tradeoff

- 2.3.1 Critical path
- 2.3.2 Resource tradeoffs
- 2.3.3 End-to-end vs. hop-by-hop
- 2.3.4 Protocol layering
- 2.3.5 State and Hierarchy
- 2.3.6 Control mechanisms
- 2.3.7 Distribution of application date
- 2.3.8 Protocol data units

2.4 Design techniques

Protocol Layering

"Layering is a useful abstraction for thinking about networking system architecture and for organizing protocols based on network structure."

- Layering is a useful abstraction – role-based

1 and 2	Link
3 and 4	Routers
4, 5, 6, and 7	End system

- Layering provides service abstraction – isolates protocols

Protocol Layering

- Effects
  1. Layering poor implementation technique
  2. Redundant layer functionality
  3. Layer synergy
  4. Hourglass
  5. Integrated layer processing
  6. Balance transparency and abstraction vs. hiding
  7. Support a variety of interface mechanisms
  8. Interrupt vs. polling
  9. Interface scalability

Protocol Layering

- Layering poor implementation technique
  - "Layering protocol architecture should not be confused with inefficient layer implementation techniques."
  - Specifies service abstraction
  - Doesn't specify implementation

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
  - "Functionality should not be included at a layer that must be duplicated at a higher layer, unless there is a performance benefit in doing so."
  - Functionality should not be included in a layer
    - That must be located at a higher layer (E2E argument)
    - Unless an overall performance benefit (HBH effect)
  - E2E vs. A2A (application-to-application) functionality

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
  - "When layering is used as a means of protocol division, allowing asynchronous protocol processing and independent data encapsulations, the processing and control mechanisms should not interfere with one another. Protocol data units should translate efficiently between layers.

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
- Hourglass
  - "The network layer provides the convergence of addressing, routing, and signalling that ties the global Internet together."
  - Intelligent pro-active networking
    - Reduces constraints

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
- Hourglass
- Integrated Layer Processing
  - "When multiple layers are generated or terminated in a single component, all encapsulations/decapsulations should be done at once, if possible in hardware."

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
- Hourglass
- Integrated Layer Processing
- Transparency and Hiding
  - "Layering is design around abstraction, providing a simpler representation of a complicated interface."
    - Translucency is better then transparency

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
- Hourglass
- Integrated Layer Processing
- Transparency and Hiding
- Variety of Interface Mechanisms
  - "A range of interlayer interface mechanisms should be provided as appropriate for performance optimization; synchronous and asynchronous, as well as interrupt-driven and polled."

Protocol Layering

- Layering poor implementation technique
- Redundant Functionality
- Layer Synergy
- Hourglass
- Integrated Layer Processing
- Transparency and Hiding
- Variety of Interface Mechanisms
- Interrupt vs. Polled

Fundamental Reasoning

Systemic Optimization

“Networks are systems of systems with complex compositions and interactions at multiple levels of hardware and software. These pieces must be analyzed and optimized in concert with one another.”

- Effect
  1. Granularity
  2. Consider side effects
  3. Keep it simple and open
  4. System partitioning
  5. Flexibility and workaround

Systemic Optimization

- Granularity
  - "Like sand on a sea shore."
  - Methodical helps isolate concerns and provides uniform communication

Systemic Optimization

- Granularity
- Side Effects
  - "Optimizing frequently have unintended side effects to the detriment of overall performance."
  - May reduce overall performance
  - Have a long-term outlook

Systemic Optimization

- Granularity
- Side Effects
- Simple and Open
  - Difficult to understand and optimize complex system of systems

Systemic Optimization

- Granularity
- Side Effects
- Simple and Open
- System Partitioning
  - "Carefully determine how functionality is distributed across a network."
  - Analyze partitioning of functionality
    - Across enterprise
    - Among components

Systemic Optimization

- Granularity
- Side Effects
- Simple and Open
- System Partitioning
- Flexibility and Workaround
  - “Provide protocol fields, control mechanisms, and software hooks to allow graceful enhancements and workarounds when fundamental tradeoffs change.”

Systemic Optimization

- Granularity
- Side Effects
- Simple and Open
- System Partitioning
- Flexibility and Workaround
- Second-Order Effects
  - “Components with only a second-order effect on performance should not be the target of optimization.”
    - i.e. optimizing a link that is not a bottleneck
    - i.e. optimizing a LAN latency for a WAN application
    - i.e. optimizing operations not in critical path

Table of Contents

Network Architecture

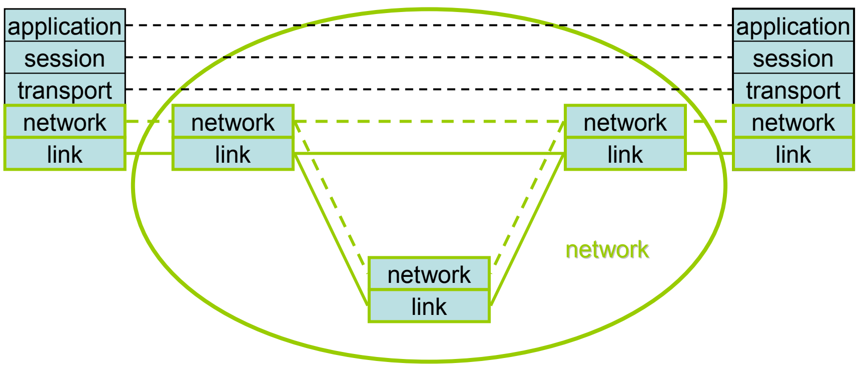

3.1 Topology and geography
3.2 Scale
3.3 Resource Tradeoffs

Network Architecture

Scalable Topologies

"*Mesh network technologies scale better than shared medium.*"

Network Architecture

Latency

"The latency along a path is the sum of all its component delays. The benefit of optimizing an individual link is directly proportional to its relative contribution to the end-to-end latency."

- $D = \sum d_i$

Latency

"The number of per hop latencies along a path is bounded by the diameter of the network. The network topology should keep the diameter low."

- *Network Latency $D = (h-1)d_s + \sum_h d_i$*
  - *d_i = component delays and bounded by speed of light*
  - *h = number of hops*
  - *d_s = routing delay*

Network Architecture

3.1 Topology
- 3.1.1 Scalability
- 3.1.2 Latency
- 3.1.3 Bandwidth
- 3.1.4 Virtual overlays and lightpaths
- 3.1.5 Practical constraints

3.2 Scale

3.3 Resource Tradeoffs

Bandwidth

"The maximum bandwidth along a path is limited by the minimum bandwidth link or component."

Network Architecture

3.1 Topology
- 3.1.1 Scalability
- 3.1.2 Latency
- 3.1.3 Bandwidth
- 3.1.4 Virtual overlays and lightpaths
- 3.1.5 Practical constraints

3.2 Scale

3.3 Resource Tradeoffs

Overlay Networks

"Overlay networks must provide the same high-performance paths as the physical networks. Overlays must be adaptable based on end-to-end path requirements."

- Layer 2 VPNs
- Layer 3 VPNs
- MPLS VPNs
- Tunnels
- Lightpaths

Network Architecture

3.1 Topology

3.2 Scale

- 3.2.1 Network Engineering
- 3.2.2 Hierarchy
- 3.2.3 Bandwidth aggregation and isolation
- 3.2.4 Latency optimization
- 3.2.5 Practical constraints

3.3 Resource Tradeoffs

Network Engineering

"Networks should be able to scale in size while balancing switch cost against hop count."

Topology	Degree	# Nodes	Diameter	Aggregation
Sparsely connected	Low	High	High	Low
Densely connected	High	Low	Low	High

Network Architecture

Hierarchy

"Use hierarchy and clustering to manage network scale and complexity, and reduce the overhead of routing algorithms."

- Control latency and aggregation
- Boundary state change
- Divide network into clusters to boundary state and aggregation

Network Architecture

3.1 Topology

3.2 Scale

- 3.2.1 Network Engineering
- 3.2.2 Hierarchy
- 3.2.3 Bandwidth aggregation and isolation
- 3.2.4 Latency optimization
- 3.2.5 Practical constraints

3.3 Resource Tradeoffs

Aggregation and Isolation

"Use hierarchy to manage bandwidth aggregation in the core of the network, and to isolate clusters of traffic locality from one another."

- Isolate bandwidth in different subnet layers

Network Architecture

Latency Optimization

"Use hierarchy and cluster size to help reduce network diameter and resultant latency."

- Limit number of hops and control aggregation

Network Architecture

Practical Constraints

"Administrative constraints increase the importance of good design."

- Policy-based routing
- External providers network

Network Architecture

3.1 Topology

3.2 Scale

3.3 Resource Tradeoffs

- 3.3.1 Bandwidth, processing and memory
- 3.3.2 Latency as a constraint
- 3.3.3 Relative scaling with high speed
- 3.3.4 Active networking

Bandwidth, Processing and Memory

"Networks are a collection of resources. Their relative composition of these resources must be balanced to optimize cost and performance and to determine network topology, engineering, and functional placement."

- Resources:
  - P = processing
  - M = memory
  - B = bandwidth
  - E = energy or power
- All constrained by:
  - L = latency

Network Architecture

3.1 Topology

3.2 Scale

3.3 Resource Tradeoffs

- 3.3.1 Bandwidth, processing and memory
- 3.3.2 Latency as a constraint
- 3.3.3 Relative scaling with high speed

Latency as a Constraint

"If all resources/constraints scale uniformly then latency become the inherit constraint."

Network Architecture

Relative Scaling

"The relative cost of resources and constraints changes over time due to non-uniform advances in different aspects of technology."

- Speed of light remains constant, so
- And Moore's law increases in P

Table of Contents

Network Architecture

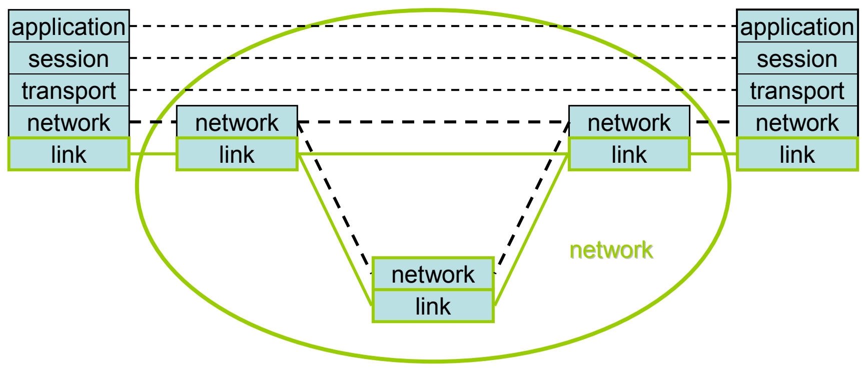

5.1.1 Physical transmission
5.1.2 Link technologies
5.1.3 Link layer components
5.1.4 Support for higher layers

Link Principle

"Network links must provide high-bandwidth connections between network nodes. Link-layer protocol processing should not introduce significant latency."

Network Components

Physical Transmission

- Physical properties
  - Signal propagation velocity
  - Link length
  - Rate
  - Encoding
  - Modulation
- High-speed network characteristics
  - Rate ≠ latency
  - Bandwidth = data rate (not spectrum bandwidth)
  - Bandwidth-x-delay product

Type	Medium	Frequency	Velocity	Delay	Attenuation
Fiber	Glass	120-250 THz	0.68c	5 µs/km	0.2-0.5 dB/km

Physical Transmission

- Fiber Optics
  - Bandwidth ≈ 20 THz within 800-1700 nm
  - Attenuation
    - Material absorption
    - Rayleigh scattering: varies by index of refraction
    - Waveguide geometric imperfections
  - Dispersion
    - Intermodal: differing distance per reflection mode
    - Chromatic: differing velocity per λ
    - polarization: differing velocity per polarized state

Physical Transmission

- Fiber Optics
  - Bandwidth ≈ 20 THz within 800-1700 nm
  - Attenuation
    - Material absorption
    - Rayleigh scattering: varies by index of refraction
    - Waveguide geometric imperfections
  - Dispersion
    - Intermodal: differing distance per reflection mode
    - Chromatic: differing velocity per λ
    - polarization: differing velocity per polarized state
  - Types
    - Single mode: single reflection mode
      - 1310 nm is white light
      - 1530 – 1610 nm is colored light

Network Components

5.1 Links

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

Link Layer Technologies

“Link protocols should be scalable in bandwidth, either variably or in discrete intervals.”

- Sequence of frame packets along a link
- Data and management planes
  - OTN: optical transport network
  - SDH: synchronous digital hierarchy
  - Ethernet
- Control Plane
  - GMPLS
  - PNNI

Link Layer Technologies

- Optical WDM
  - Number of wavelengths inversely proportional to distance
    - Stimulated Raman scattering due to molecular vibrations
    - Stimulated Brillouin scattering interacting with acoustic waves
    - Carrier-induced cross-phase modulation causes phase shifts
    - Four-wave mixing induces sum and difference frequencies
  - Contraints
    - When and where to conduct OEO

Network Components

5.1 Links

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

Link Layer Components

"Link layer components must sustain the data rate while not introducing significant delay."

- Amplifiers and regenerators
- Multiplexers and cross-connects
- Optical wavelength converters (OEO)
- Routers and switches

Network Components

5.1 Links

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

Support for Higher Layers

"Provide long-term and dynamic information on the reason for loss to high layers so that network and end-to-end mechanisms respond appropriately."

- Response depends on type of L2 and sublayer 3 loss
  - Net configuration/rerouting based on path characteristics
  - Transport layer response to corruption
- Filter PDUs not destined for higher layer protocols
- Discard early at low layers
- Infrastructure protocols need broadcast support

Impact of High Speed

"Optimize header control field values to trade efficiency against expected future requirements."

- Bandwidth-x-delay product increases
  - *Fixed duration error event larger relative impact*
  - *Reaction to errors decrease in terms of number of bits sent*
- *Sequence numbers wrap if sequence space is too small causing packet misinsertion*