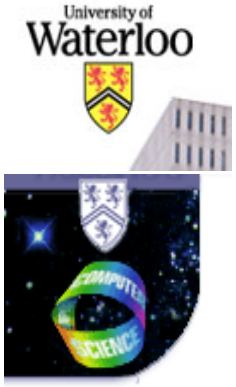

CS 856
Internet Transport Performance

Introduction

Martin Karsten

School of Computer Science, University of Waterloo

mkarsten@uwaterloo.ca

Contents

Communication Networks

Transport infrastructure for intangible goods: Information

Digital Communication Network
- **information represented as bits**
- **and aggregations thereof**

Infrastructure
- **devices, cables, antennas, laser, ...**

Services & Applications ⇒ Layering
- **e.g. ISDN 64 kbit/s channel → service**
  - call setup & speech encoding/decoding → application (telephone)
  - analog fax over resulting audio channel → application
    - "telephone application" → service

Purpose of Communication Networks
- **general vs. special**

The "Internet"

Internet (IP-based) Technology
- protocol suite: IP, DHCP, TCP, UDP, OSPF, BGP, HTTP, etc.
- specifications, software, hardware, policies

Network Domain
- collection of hosts and links
- common network technology

Administrative Domain
- common administrative authority
- coherent policy

Public Internet
- collection of network/administrative domains
- connected through least common denominator → IP
- loosely coherent global policies
  - usage of BGP routing vs. local configuration of BGP routing

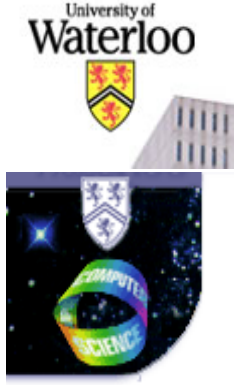

Internet Protocol Suite - Recap

ISO Open Systems Interconnection (OSI) Reference Model

7	Application Layer
6	Presentation Layer
5	Session Layer
4	Transport Layer
3	Network Layer
2	Data Link Layer
1	Physical Layer

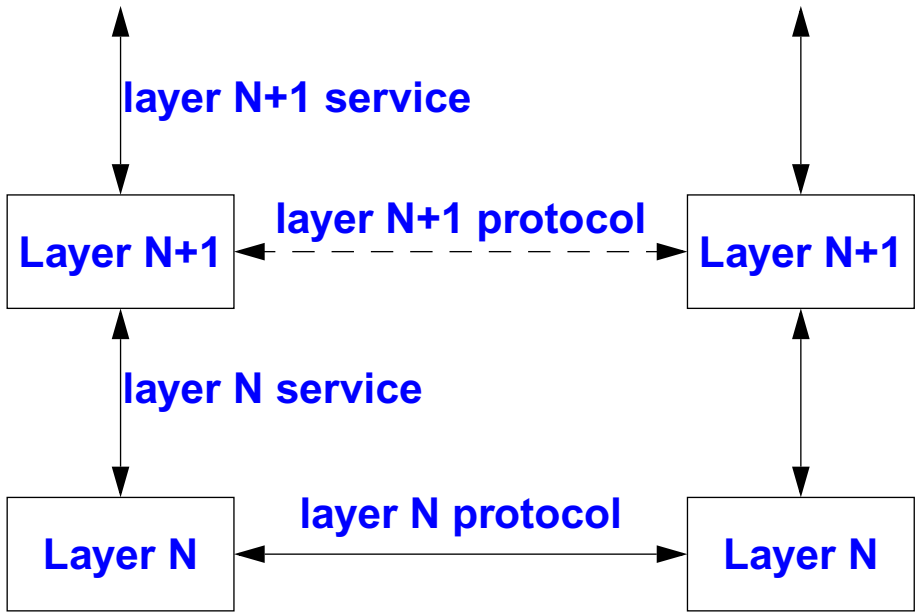

Internet Model

- layer 1-4 & application layer
- expanding?

Internet Protocol Suite - Recap

Communication Layers between Processes

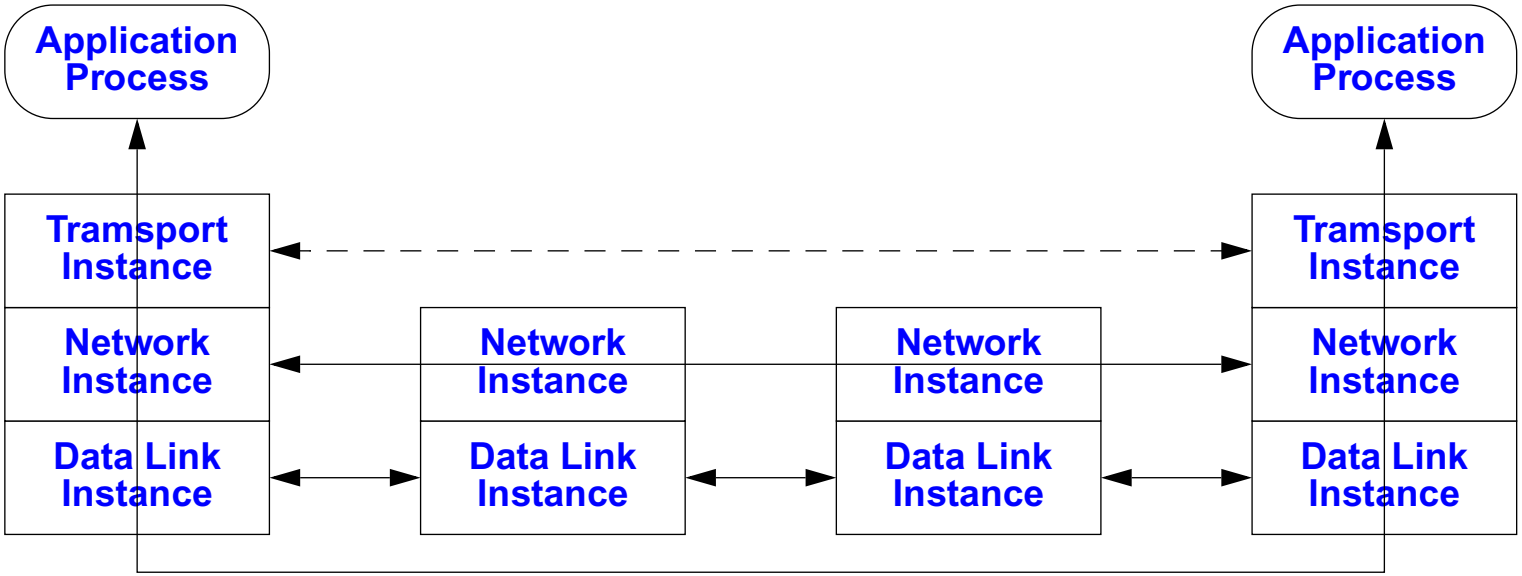

Internetworking vs. Broadcast Networks

- limited scalability of broadcast networks → LAN or MAN

Data Link Instance: Hardware/Device

- modern network devices: some TCP/IP end system functionality

Transport & Network Instance: Software

- modern routers: network instance (partially) = special hardware

Internet - Today and Tomorrow

Current Services

- remote resource access: printing, file server, remote login, etc.
- information distribution: web, file transfer, etc.
- communication: email, news, chat, voice, video

Reliability and Performance - Public Internet

- often good enough to very good, but: would you trust your life on it?
- limited regulation and no minimum guarantees
- private IP networks $\rightarrow$ different story

Vision

- uniform integrated/layered infrastructure
  - cost reduction: economies of scale
- diverse applications
- reliable services and predictable performance
- uniform public communication service?

IETF - Standardization & Research

Relevant Interfaces

- **between modules**
- **between devices**
- **between networks**

Internet Engineering Task Force

- **standardization of protocols "on the wire"**
- **operates by "consensus" rather than formal voting**
  - group of individuals rather than organizations
  - no formal membership → open meetings and debate
- **"running code" principle → existing inter-operable systems required**

IETF ↔ Research Community

- **(decreasing) overlap**
- **(increasing) influence of vendors**
- **experimental documents (Internet Drafts, RFCs)**

Note: Standardization → Technology AND Marketing !

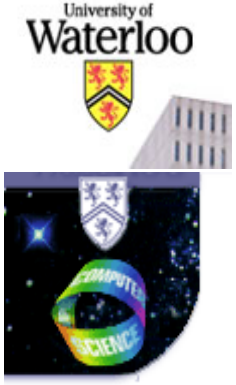

Topics - Network Control

Routing

- establish connectivity between end systems

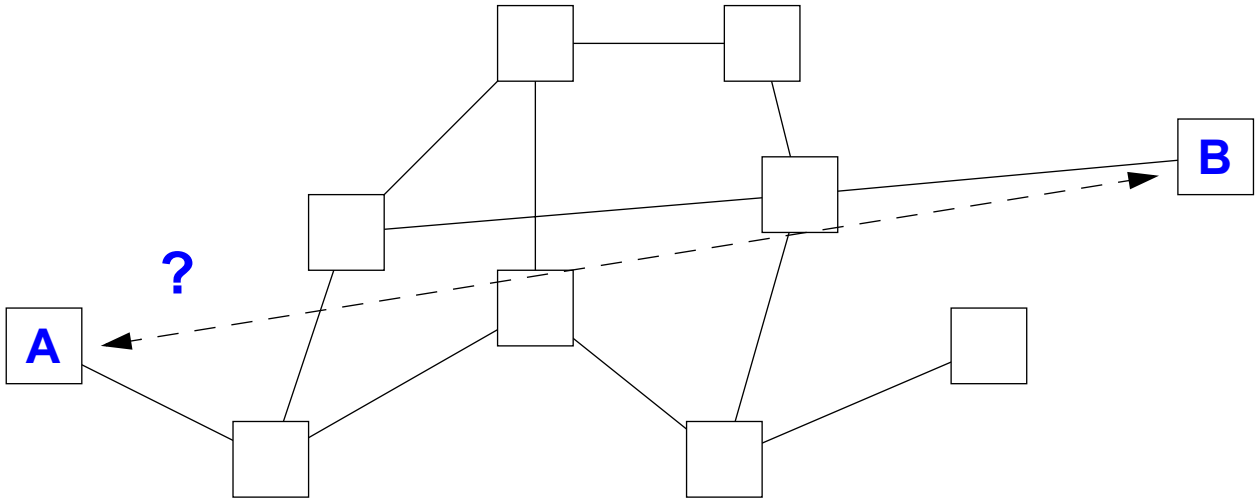

- does at least one path exist?
- if multiple, which one to take?
- symmetric vs. asymmetric routing?

Routing Mechanisms

- information gathering: global vs. local vs. centralized
- path establishment: implicit vs. explicit

Topics - Network Control

Connectivity & Multiplexing - Taxonomy

- **circuit switching → no multiplexing**
  - establishment of closed physical end-to-end path
    - old telephone networks → electrical circuit
    - fiber-optics networks (WDM) → optical circuit
- **time division multiplexing → synchronous multiplexing**
  - time slot ↔ fixed-size basic information unit (X bits per cell)
  - synchronous transmission, path identification through slot number
  - explicit path setup required
    - modern/current telephone networks (Sonet/SDH)
- **virtual circuit → asynchronous multiplexing with explicit path setup**
  - fixed- or variable-sized information unit
  - asynchronous transmission, path identification through label/identifier
  - label/identifier information exchanged during path setup
    - ATM (asynchronous transfer mode)
- **datagram → asynchronous multiplexing without explicit path setup**
  - fixed- or variable-sized information unit
  - asynchronous transmission, path identification through end system address
  - compact addressing required
  - independent routing function required
    - Internet

Topics - Network Control

Routing & Aggregation

- **global scope: millions of end systems → routing state?**
- **aggregation necessary**
  - simple aggregation: e.g. class A,B,C subnets in IPv4 routing
- **with and without explicit path setup**

Internet Routing

- **intra-domain → efficiency (performance and load control)**
  - link-state routing provides global information
  - requires increased communication and computation
- **inter-domain → policy and economics**
  - distance-/path-vector routing → route advertising
  - local optimization vs. global consistency?

Inter-Domain Internet Routing

- **old problem: flexibility → classless inter-domain routing (CIDR)**
  - requires more complicated lookup function (longest-prefix matching)
- **new problem: exponential growth of routing state**
  - multi-homing: multiple advertisements for each network

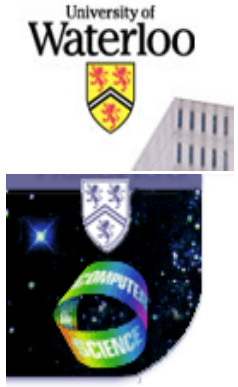

Topics - Network Control

Network Signalling

- **establish connectivity**
  - set up virtual circuit
  - find mobile end system
  - open firewall
  - etc.
- **negotiate traffic and service contract**
  - QoS signalling: interface vs. distributed algorithm
- **overhead of network signalling?**

Application-level Signalling

- **e.g. initiate voice call**
- **private networks: integrate application & network signalling**

Failure Handling

- **node link - overload and/or failure**
- **detect problematic condition and reroute traffic**
- **information gathering → execute necessary changes**

Topics - Packet Scheduling

Contention

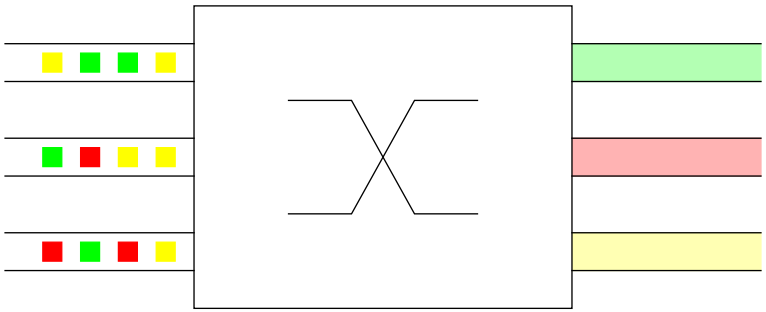

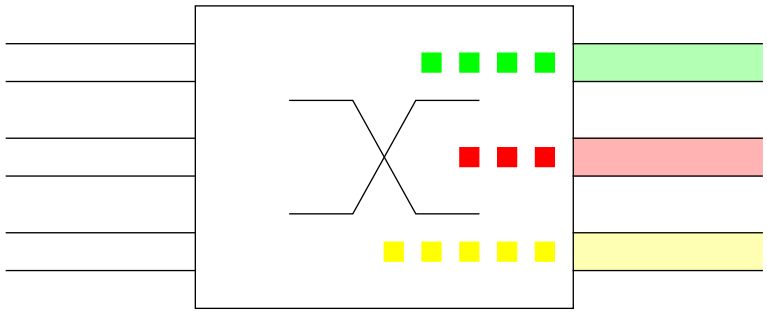

Input Queueing - Potential Bottlenecks

- classification & routing lookup
- access to backplane
- switching capacity of backplane

Output Queueing - Potential Bottlenecks

- link capacity
- link access (broadcast medium)

In Reality: Both

Topics - Packet Scheduling

Asynchronous Multiplexing: (Unregulated) Arrival at Multiple Input Ports

- buffering → delay
- buffer overflow → drop

Diverse Requirements

- throughput
- delay
- loss

Packet Scheduling

- which packet to transmit next?
- buffer management
- granularity of discrimination → application flow vs. traffic class
- relative vs. absolute service guarantees
- path setup ↔ traffic regulation & state setup

Fairness

- definition of fair resource sharing?

Topics - Congestion Control

Datagram Network and Aggregated Routing

- no explicit path setup
- no pro-active resource allocation
- transient congestion will happen

Congestion Control

- original goal: regulate fast sender by slower receiver
- network congestion: regulate collection of senders by network
- contain congestion as fast as possible
- approximate available capacity (fair share)

Effects

- implicit network control through feedback loop
- stable resource allocation $\rightarrow$ same effect as traditional network control?

Challenges

- large bandwidth/delay product and reaction delay
- high utilization vs. careful approximation of fair share of capacity

Topics - Mobility & Multicast

Advanced Communication Patterns

- group communication
- (m)anycast communication
- end system mobility
- middle-box functionality (e.g. firewalls)

⇒ **Not well handled by connectionless datagram service**

- e.g. mobility: information about current network access point needed
- e.g. multicast: information about current group membership needed

Existing Internet Proposals

- consider advanced communication patterns as exception
- state-based additions to IP (end system based → overlay networks)
- dual-stack, specialized solutions
- no changes to underlying IP operation

What if these "exceptions" become the norm?

Topics - Network Architecture

Current Internet Architecture

- **unicast data communication**
  - human-to-human communication?
  - real-time communication?
  - group communication?
- **earth-based, fixed, wired workstations**
  - variety of devices?
  - variety of access and transit links?
  - device mobility?
- **basic services**
  - extremely high reliability requirements?
  - extremely security requirements?

Case for Integrated Infrastructure

- **high efficiency needed**
- **engineering for the common case**
- **common case → moving target?**

Implementation of Network and Transport Functionality

Goals: High Efficiency, Low Complexity, High Flexibility, High Robustness

- **often conflicting goals**

Conceptual Levels

- **data plane - everything directly involved in moving packets**
- **control plane - configuration of data plane: routing, signalling, etc.**

Implementation Levels

- **software**
  - user-level processes
  - kernel-level modules
- **hardware**
  - various programming models

System Types

- **end system: workstations, servers, thin clients, etc.**
- **router: location in network $\leftrightarrow$ expected traffic**
- **higher-level gateway: firewall, VoIP call server, etc.**

Assessment of Conceptions

Scope

- **node, network domain, or global Internet?**
- **special vs. general purpose**
- **isolated vs. universal function**
- **anticipated application & traffic mix**
- **scalability**

Performance vs. Control

- **average performance vs. worst-case performance**
  - 90% argument vs. 99.99% argument
- **explicit control & configuration**
- **robustness & reliability → handling of failures**
- **economics & value-chain**
- **security implications**

Course Projects

Goals

- **thoroughly understand problem domain**
- **demonstrate understanding by project results and paper**
  - design sensible mechanisms
  - meaningful performance investigations
  - sketch path to "perfect" solution
  - or: demonstrate that you have explored a large variety of alternatives
- **practical experience in simulation and/or implementation and testing**
- **do more than just literature research**
- **unusual ideas are preferred (but hard to produce)**

Course Project

- **1-3 students per project → project sizing must be appropriate**
- **pick from list or propose your own!**

Project Proposal

- **one page about problem and proposed work**
- **required infrastructure**
- **discuss with instructor before submitting proposal**