Communications Network Design

lecture 04

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

March 2, 2009

Network Optimization: Goals and Constraints

What are the typical optimization goals (e.g., cost, performance, reliability) for network operators? Where are the costs in networks? What are the constraints (technological, and non-tech.) they operate under?

The lecture describes some common network optimization goals and constraints, and why each occur.

Lecture goals/outline

- Understand what optimization means
  - optimization goals
    - e.g. reduce cost
    - e.g. improve cost or reliability
  - optimization constraints
    - technological, geographic, political, ...
- think about these in a real context
  - e.g. what are the costs?
    - e.g. what is a router
  - e.g. what data do we need?
- references: for more details on Routers see Packet Switch Architectures - I, N. McKeown, B. Prabhakar
  http://www.stanford.edu/class/ee384x/syllabus.html

Network Optimization Goals

- costs (usually assume equipment costs are large)
- performance (minimize delays, or latency)
- survivability
  - hard to write as an optimization problem
  - heuristic approach
    - distributed network
    - redundancy

Communications Network Design: lecture 04 – p.3/39

Communications Network Design: lecture 04 – p.4/39

Cost in networking

Arguments about which costs are biggest

- capital
  - ▷ equipment (cables, switches, ...)
  - ▷ premises
    - ⋆ + land that cables run along (right of ways)
- operations
  - ▷ **exclude** sales and marketing, management, R&D
    - ⋆ doesn't depend on network design
  - ▷ salaries of network administrators
    - ⋆ repairs and upgrades
    - ⋆ **design**
  - ▷ power
  - ▷ transit (from upstream providers)
    - ⋆ fixed
    - ⋆ traffic based costs

Equipment costs

Often assumed to dominate

- fixed node costs
  - ▷ cost of a router - often assumed small
  - ▷ need to include premises, installation, etc.
- fixed link costs
  - ▷ constant component
  - ▷ BW component
    - ⋆ higher bandwidth links cost more
- distance costs
  - ▷ straight distance cost
  - ▷ BW x distance cost

Telstra half-year report, 2005, [1]

- revenue >$11 billion
- profit >$2 billion
- property, plant, equipment >$22 billion

So in this case, operational costs per year are almost half the total capital value for the network. However, that's a bit misleading, as a lot of the costs here aren't networking costs: they are the costs of sales and marketing, management salaries, billing, legal costs, research and development, and so on. It's often a bit hard to disentagle the operational costs that you (as network designer) have control over from "administration" costs.

Communications Network Design: lecture 04 – p.5/39

Communications Network Design: lecture 04 – p.6/39

Link costs

Linear model: cost of a link

$$\text{Cost} = k + \alpha r + \beta d + \gamma r d$$

where

$$\begin{aligned} r &= \text{link capacity} \\ d &= \text{link distance} \end{aligned}$$

- the parameters k, α, β, γ are constants.
- often some terms might be close to zero so ignore
- some terms are out of our control, so we ignore these, or push them into constants

Simple Example Problem

Lets consider the problem of business that wants to connect up two locations with a 10 Mbps link. What can they do:

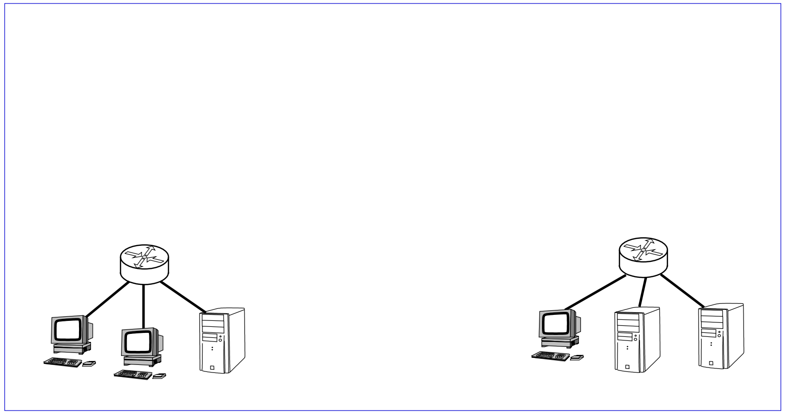

Simple Example Problem

We have two possible solutions:

- private line
  - lease or build whole line
  - cost depends on distance: $C = k_{\text{private}} + \beta_{\text{private}} d$
- VPN
  - pay for access to network at each end, but not for the network
  - no distance dependence: $\beta_{\text{VPN}} \simeq 0$
- decision: use private line if

$$k_{\text{private}} + \beta_{\text{private}} d \leq 2k_{\text{VPN}}$$

In the above problem we require a fixed capacity. As this is not variable, we ignore the capacity related cost terms, or add them into the constants.

Obviously this is a simplified example. Just designed to show something of how cost model effects problems, and how it relates to underlying tech.

The "constants"

Assume the linear model, how would you work out k, α, β, γ

- β and γ arise from the costs of building a links.
  - β are the fixed costs: right-of-way, digging cables in, i.e., things we need regardless of how much capacity we use.
  - γ reflects capacity related costs: e.g., in the old days, if you wanted two links, you needed two cables. Today, this might reflect the number of λ (wavelengths) you use on a WDM system.
- in reality, we often purchase such links from a physical layer network provider. They pass on a range of their costs through a pricing model that determines β and γ.

Remember in the linear model

$$\text{Cost} = k + \alpha r + \beta d + \gamma r d$$

The "constants"

Assume the linear model, how would you work out k, α, β, γ

- α and k represent the non-distance dependent costs of a link. These are usually associated with end equipment, for instance the WDM multiplexers, and line cards at the routers that terminate the link:
  - k is non-capacity dependent costs: cost of getting someone to install a line card, and spend time configuring the router.
  - α is capacity related term: higher speed line cards usually cost more.

To understand some of this terminology we have to understand more about what a router is.

What is a router?

A Juniper router in use.

Remember in the linear model

$$\text{Cost} = k + \alpha r + \beta d + \gamma r d$$

Logical Router

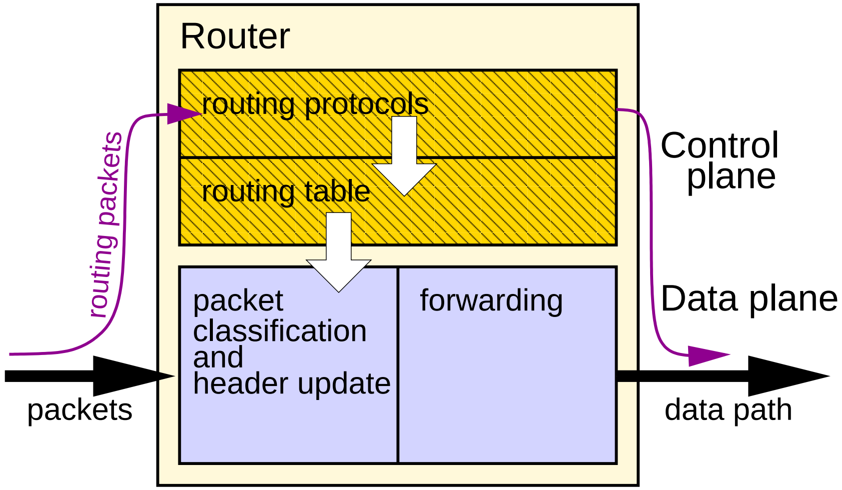

Router Architecture

Common modern architecture

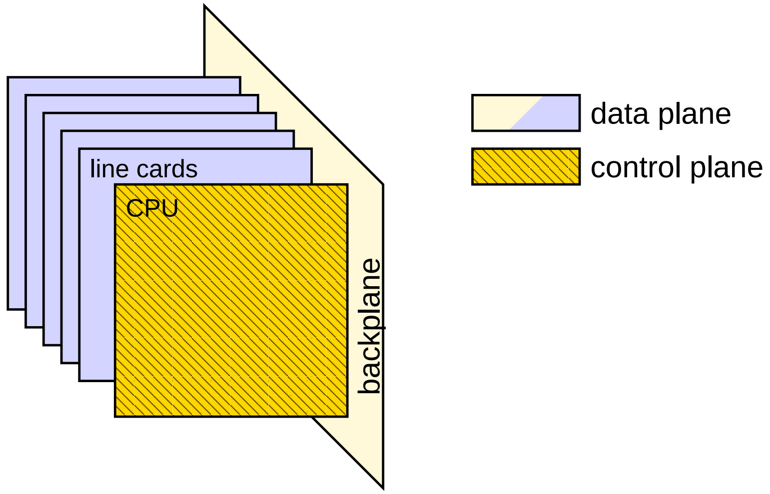

Communications Network Design: lecture 04 – p.13/39

Line card

Procket line card

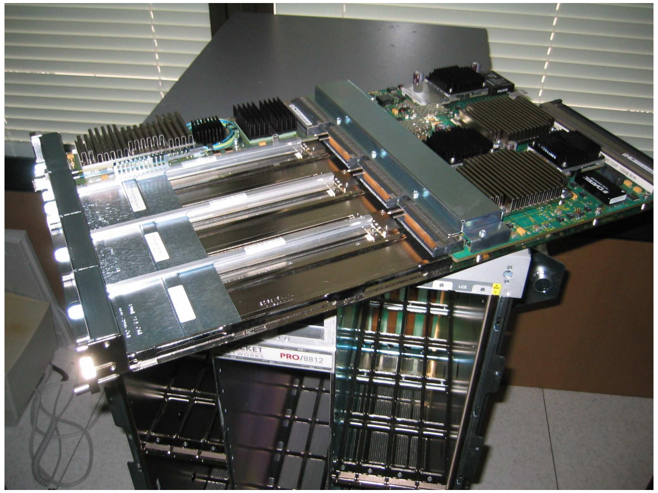

Courtesy of AARNET

CPU

Procket CPU

Courtesy of AARNET

Chassis

Procket Chassis

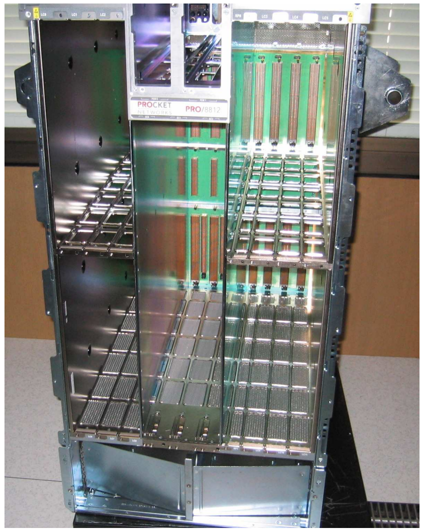

Courtesy of AARNET

Router Architecture

Less efficient software router

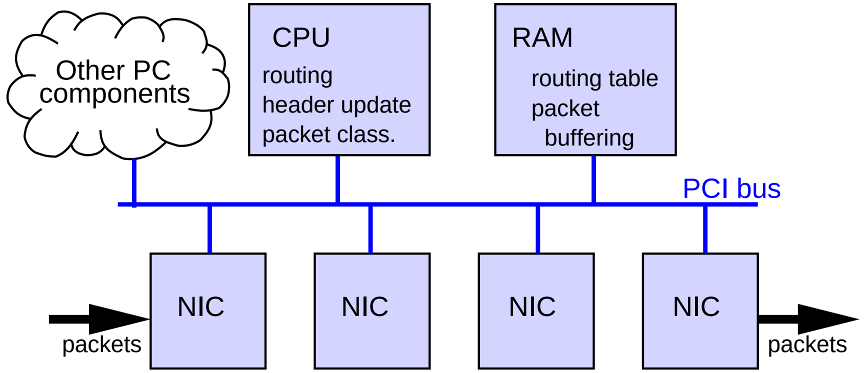

NIC = Network Interface Card

Router Architecture

High perf. architecture (input and output queueing)

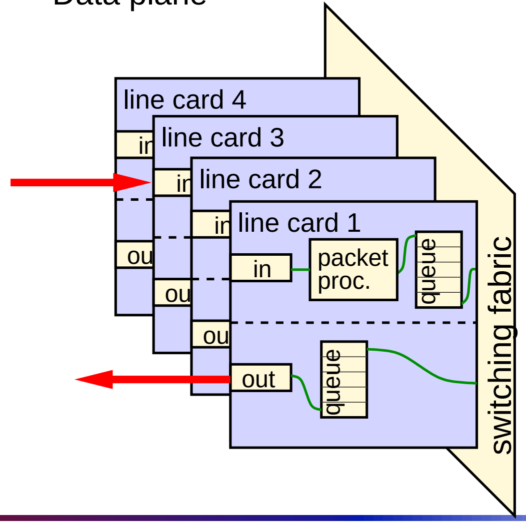

Router Architecture

High perf. architecture (input and output queueing)

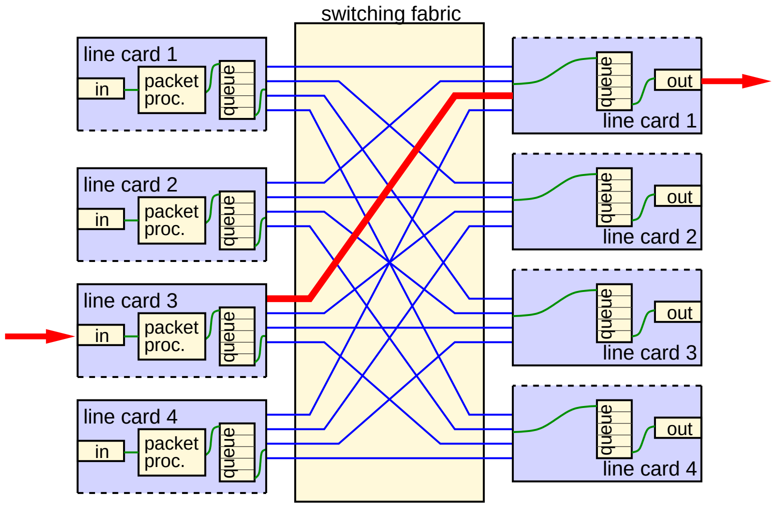

Per packet processing

In an IP Router

- lookup packet destination in forwarding table
  - up to 150,000 entries (today)
- update header (e.g. checksum, and TTL)
- send packet to outgoing port
- buffer packet along the way

For a 10 Gbps line

- small 40 byte packets
- about 30 million packets per second
- you have ∼30ns per packet

BGP routing table size

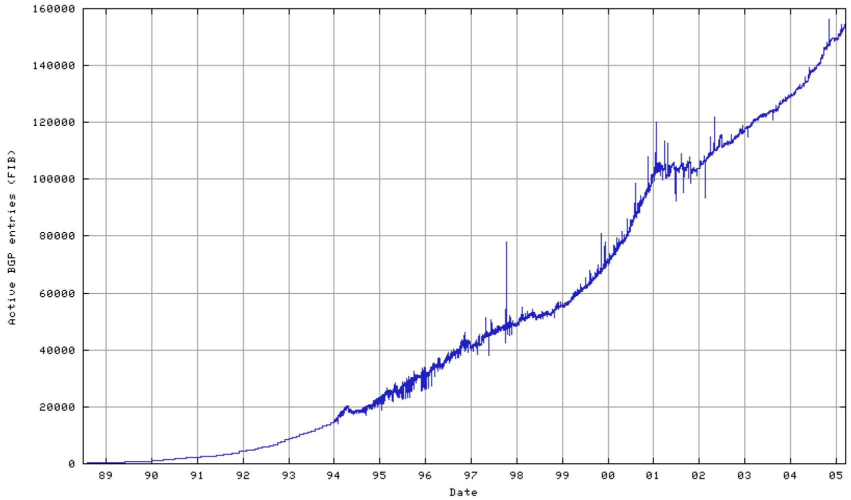

http://www.cidr-report.org/

Communications Network Design: lecture 04 – p.21/39

Communications Network Design: lecture 04 – p.22/39

Expensive bits

- forwarding table can be large
  - up to 150,000 entries per line card
  - lookup in ∼30ns for 10 Gbps line
  - need fast memory
- buffers can be large
  - 0.2 seconds per line card (rule of thumb)
  - 10 Gbps line = 250 MB memory (on in and out)
  - need fast memory (in + out in ∼30ns)
- backplane must be faster than line cards
  - N times line rate speedup (N linecards)
  - to guarantee non-blocking switch fabric

Router costs

- chassis
  - one time cost per router
  - but depends which chassis
  - large (more expensive) chassis fits more line cards
- line card
  - number of ports
  - speed of ports
  - Cisco 12000 Series examples
    - ⋆ Eight-Port Fast Ethernet Line Card
    - ⋆ Router Gigabit Ethernet Line Card
    - ⋆ Three-Port Gigabit Ethernet Line Card
    - ⋆ 10-Port Gigabit Ethernet Line Card

Link costs alternatives

- distance component of physical link
  - wired: cost of fibre, amplifiers/repeaters, digging, right of way
  - wireless: (e.g., free-space optics) free over short distances
- logical link (VPN-like networks)
  - (simplified) cost depend on capacity, but not distance
  - may depend on actual traffic volume
- satellites
- big companies often vertically integrated
  - internal sales of bandwidth between divisions

Linear model: what's it good for?

- is a linear model of costs good?
  - not really
- in terms of costs, this is a **discrete** problem
  - but its too complicated
  - hard to get exact pricing info anyway
    - pricing often depends on size of order, or internal company politics
- we will often treat it as linear (continuous)
  - as an approximation
  - note that a major source of inefficiency is in the discrete nature of bandwidths, and router capabilities

Communications Network Design: lecture 04 – p.25/39

Real link costs aren't linear. Links come in discrete capacities. Common link speeds.

- T1: 1.544Mbps
- E1: 2.048Mbps
- T3: 44.736Mbps (=28xT1)
- DS3: 44.736Mbps
- OC3/STM1: 155.52 Mbps (=100 T1)
- OC12/STM4: 622.08 Mbps (=4xOC3)
- OC48/STM16: 2.488 Gbps (=4xOC12)
- OC192/STM64: 9.953 Gbps (=4xOC48)
- OC768/STM256: 39.813 Gbps (=4xOC192)
- Ethernet (10BaseT): 10 Mbps
- Fast-Ethernet: 100 Mbps
- Gig-E: 1 Gbps
- 10Gig-E: 10 Gbps

Pricing is non-linear, depending on what vendors want to push today. Vendors also often give discounts for multiple orders.

Communications Network Design: lecture 04 – p.26/39

Optimizing for Latency

Another goal for optimization is to maximize network performance.

- network performance often measured by **latency**
- **latency** is the delay of a packet crossing the network
- most often we are concerned with average latency
  - over all paths through the network

Communications Network Design: lecture 04 – p.27/39

Optimizing for Latency

Types of delay

- **propagation:**
  - propagation delay directly related to distance
- **queueing:**
  - queueing is caused by transient congestion
- **processing:**
  - packet processing time (address lookup, and header update)
  - fixed per hop
- **transmission:**
  - time to tranmit packet on the line
    = packet size / line rate

Communications Network Design: lecture 04 – p.28/39

Different scenarios

- ARPANET low speed links (56 kbps), and slow processors (IMPs)
  - **propagation**: coast-to-coast in US ~ 30ms
  - **transmission**: $1500 \times 8/56000 = 0.22$ seconds.
  - **queueing**: a couple of packets $\sim$ a few seconds
  - **processing**: similar order to trans, but smaller.

  so transmission and queueing times dominate.
- modern national backbone (10 Gbps)
  - **propagation**: coast-to-coast in US ~ 30ms
  - **transmission**: $1500 \times 8/1.0e10 = 1.2$ ns.
  - **queueing**: large buffers (up to 0.2 seconds)
  - **processing**: ~ 30 ns.

  so queueing is dominant, unless low load, where propagation becomes dominant.

Optimizing for Latency

How to reduce

- **propagation**:
  - cannot speed up light
  - really minimizing length of paths
- **queueing**:
  - reduce queueing by reducing load
- **processing**:
  - minimizing number of hops
- **transmission**:
  - minimizing packet sizes
    - e.g. VoIP uses small packets

Optimizing for survivability

The 6 things network engineers care about

Five 9's

Goal of many telecom level providers is

- five nines reliability
- e.g. in IP networks
  - uptime is 99.999%
  - translates to about 5 minutes downtime per year
- pretty hard to achieve
  - not just network design
  - disaster recovery processes

Communications Network Design: lecture 04 – p.31/39

Communications Network Design: lecture 04 – p.32/39

Approach

Often not approached using optimization

- redundancy
  - routers, links, power supplies, A/C, ...
- distribution of control
- problem detection and diagnosis
  - network post-mortems
- disaster recovery

We will consider some optimization approaches later in the coarse (if we get time).

Technological Constraints

The other aspect of optimization is the constraints

- max node degree
  - max number of line cards per router
  - times max ports per card
- max capacity per link
  - limited by speed of line cards
  - at best follows Moore's law
  - today, around OC762 = 40 Gbps
- max capacity per router
  - backplane technology limited (also Moore's law)
  - today, around 10 Tbps
- max length of a link (e.g. Ethernet)

Maximum link distances: 185 meters for 10Base2, 500m for 10base5, and 100m for 10baseT.

Non-technological Constraints

- geography
  - cost of cable in oceans is different from land
  - expensive to lay cable in some places
    - e.g. downtown Manhattan
- politics
  - internal company organization mandates network organization
  - marketing get a better network than accounting, even though they have less real need
- security
  - may not want to share network resources outside of secure building

Other Constraints

- what if we have more than one objective
  - e.g. network should be
    - fastest
    - cheapest, and
    - most reliable
- multi-objective optimization is hard
- use other objectives as constraints, e.g.
  - best performance within a budget
  - cheapest network which meets performance constraints
  - cheapest network which meets reliability constraints

Communications Network Design: lecture 04 – p.35/39

Communications Network Design: lecture 04 – p.36/39

Other issues

- usually there are other inputs to optimization
  - traffic measurements
  - not always as easy to get as you might think
- planning horizon
  - usually when we design a network it takes some time to build
- often we can't design our network from scratch
  - have to deal with legacy equipment
  - incremental design

Network Optimization

- note we apply methods to Internet
- optimization methods are much more widely applicable
  - other networks: transport, post, air travel, ...
  - other non-network problems that can be written in the form of a network

Communications Network Design: lecture 04 – p.37/39

Communications Network Design: lecture 04 – p.38/39

References

[1] "Telstra corporation limited — half-year report," 2005.