Congestion and Congestion Control

S. Keshav
University of Waterloo

Outline

- Fundamentals
- Congestion and congestion control
- Complications
- Summary

Fundamentals

Three key concepts

- Traffic types
- Spatial and temporal multiplexing
- Strict and loose partitioning

Traffic types

- Smooth traffic
  - peak ~= average rate
- Bursty traffic
  - peak >> average rate
  - assume there is an intrinsic average rate, for now

Spatial and temporal statistical multiplexing

- Spatial statistical multiplexing (stat mux)
  - If n users of a resource, only $m < n$ are simultaneously active
    - can overbook: resource size = $m < n$
- Temporal stat mux
  - Of the set of m ($m < n$) active users, not all are transmitting at peak rate
    - can overbook: resource size $< pm$

Traffic and muxing

- For smooth traffic, only spatial muxing possible
- For bursty traffic, both spatial and temporal muxing is possible

Strict resource sharing

- Strict partitioning of networking resources at all contention points
  - by means of call setup and path pinning
    - strict isolation
    - no congestion
    - only spatial stat mux
    - no temporal stat mux

Strict: contd.

- To deal with potential overbooking with spatial stat mux, may have to deny calls
- Size network so that call blocking probability is low
- Admitted calls are isolated, and can experience no congestion

Loose resource sharing

- Allow both spatial and temporal stat muxing
- Two cases: with and without call setup
- With call setup
  - setup packet carries measure of burstiness
  - combine measures, assuming independence
  - for 'large enough' users in system, P (overload) < threshold

Loose sharing contd.

- Even with call setup, may have losses
  - because measures not combined properly
  - because sources not characterized properly
  - because of bad luck

Loose sharing without call setup

- Sources send data, not knowing whether
  - there are other sources at all (spatial stat mux)
  - or if these sources happen to be sending a burst (temporal stat. mux)

Emphasize: two problems

- Lots of sources of traffic
  - due to spatial stat mux
- Only a few sources, but all sending a burst
  - due to temporal stat mux

Two time scales

- If not too many sources, but all sending a burst (temporal)
  - bursts will dissipate ‘soon’
- If too many sources (spatial)
  - traffic will not dissipate any time soon
  - need to turn down the average rate!

Control mechanisms

- If not too many sources, but bursts happen to coincide
  - use a buffer to absorb bursts, no need to change average rates
- If too many sources
  - need to change average rates

Congestion and congestion control

Source perspective

- Congestion is experienced at a source
- If bursts coincide
  - increased queueing delay, but may not be noticeable
  - potentially packet loss
- If too many sources
  - persistent queueing
  - packet loss almost certain

Source perspective contd.

- At low loads, source is uncoupled with other sources
- At higher loads, source is coupled with other sources
  - increase in overall load causes more delay, and potentially packet loss

Congestion

- Decrease in utility, from the perspective of a source, due to increased load

Control action at a source

- Depends on the source utility function
- If its utility does not decrease due to increased delays or losses, then do nothing
- If its utility decreases, do something to recover lost utility
  - what it does depends on the incentive structure set up by the system

Network perspective

- Short term load imbalances are unavoidable
  - easily corrected by buffering
- But too much buffering can lead to persistent queueing and persistent delay
  - limit to buffering
- For long term load imbalance, need to set up an incentive structure

Incentive structure

- Must make it incentive compatible for sources to reduce their average rates
  - but this is tricky - reducing average rate almost certainly reduces utility
  - so why would a source do it?
  - and not renege, so that some other source does the reduction?

The heart of the problem

- How can we ensure that sources reduce their rates when it reduces their utility?
- How can we prevent reneging?

Solution schemas

- Legislate cooperative sources
- Forcibly reduce load in proportional to source rates (RED)
- Decouple rates, and let each source deal with its own mess
  - needs scheduling

Complications

Inherent problems

- Overhead
- Fairness
- Efficiency
- Stability
- Heterogeneity
- Legacy

Overhead

- How to tell sources about overload?
  - Explicit or implicit
- Explicit introduces overhead

Fairness

- Reduction of overload implicitly allocates the transmission and memory resources
- Should this allocation be fair?
  - What does fairness mean anyway?

Efficiency

- If reduction in load is too much, then the network is underloaded
  - loss of efficiency

Stability

- Overlead -> signal -> underload -> signal -> oveload
  - oscillating load, which can lead to inefficiency and poor utility

Heterogeneity

- Sources differ in
  - need for capacity
  - need for delay
  - ability to respond
  - end-to-end round trip time
- Operators differ in
  - operating costs
  - capital structure

Legacy

- ~1 billion sources and ~10 milliion routers already deployed!
- More than the hardware cost, is the existing wetware - professional training

Summary

- Congestion is reduction in utility due to overload
  - in networks that support both spatial and temporal multiplexing, but no reservation
- Solution schemas exist, but inherently complex