DATA AND COMPUTER COMMUNICATIONS

Lecture 4 Wide Area Networks - Congestion in Data Networks

Mei Yang

Based on Lecture slides by
William Stallings

1

What Is Congestion?

- congestion occurs when the number of packets being transmitted through the network approaches the packet handling capacity of the network
- congestion control aims to keep number of packets below a level at which performance falls off dramatically
- a data network is a network of queues
- generally 80% utilization is critical
- finite queues mean data may be lost

Queues at a Node

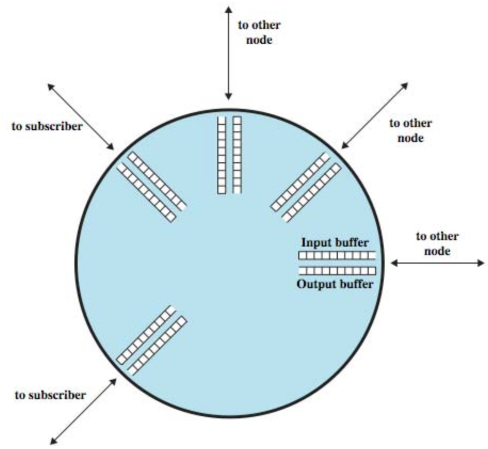

Interaction of Queues

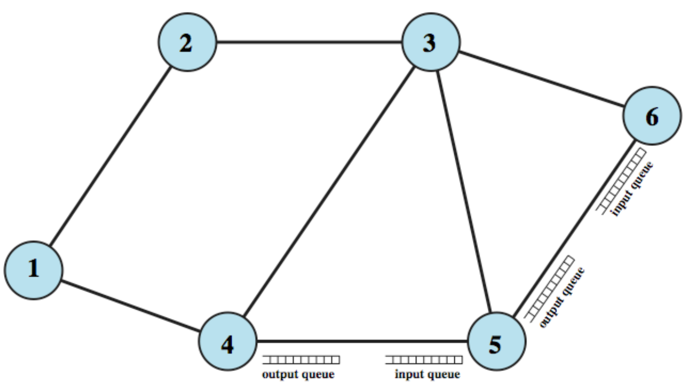

Ideal Network Utilization

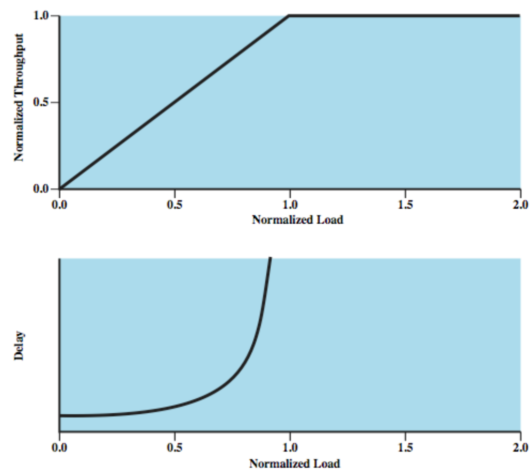

Effects of Congestion - No Control

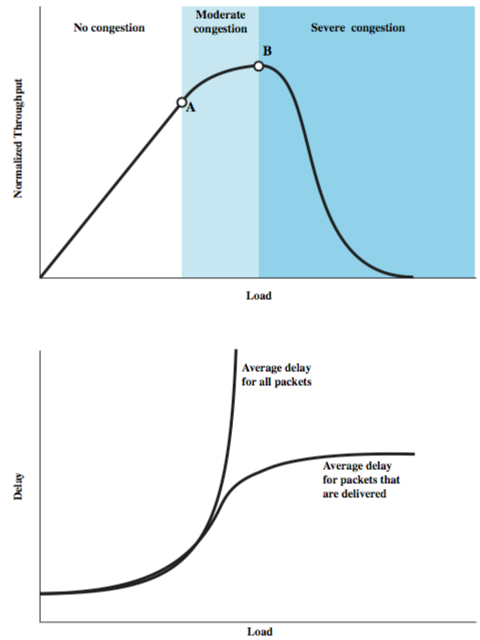

Mechanisms for Congestion Control

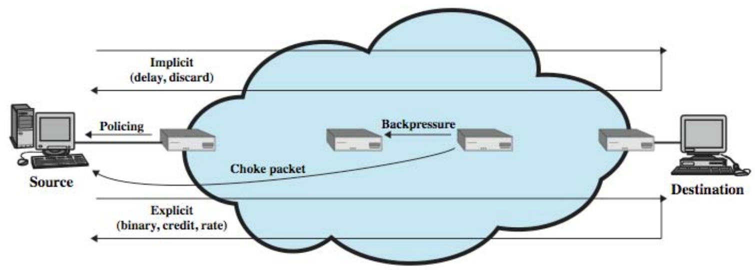

Backpressure

- if node becomes congested it can slow down or halt flow of packets from other nodes
  - cf. backpressure in blocked fluid pipe
  - may mean that other nodes have to apply control on incoming packet rates
  - propagates back to source
- can restrict to high traffic logical connections
- used in connection oriented nets that allow hop by hop congestion control (eg. X.25)
- not used in ATM nor frame relay
- only recently developed for IP

CHOKE PACKET

- a control packet
  - generated at congested node
  - sent to source node
  - eg. ICMP source quench
    - from router or destination
    - source cuts back until no more source quench message
    - sent for every discarded packet, or anticipated
- is a rather crude mechanism

IMPLICIT CONGESTION SIGNALING

- transmission delay increases with congestion
- hence a packet may be discarded
- source detects this implicit congestion indication
- useful on connectionless (datagram) networks
  - eg. IP based
    - (TCP includes congestion and flow control - see chapter 17)
- used in frame relay LAPF

Explicit Congestion Signaling

- network alerts end systems of increasing congestion
- end systems take steps to reduce offered load
- Backwards
  - congestion avoidance notification in opposite direction to packet required
- Forwards
  - congestion avoidance notification in same direction as packet required

Explicit Signaling Categories

Binary	Credit based	Rate based
• a bit set in a packet indicates congestion	• indicates how many packets source may send • common for end-to-end flow control	• supply explicit data rate limit • nodes along path may request rate reduction

Traffic Management

fairness

- provide equal treatment of various flows

quality of service

- different treatment for different connections

reservations

- traffic contract between user and network
- excess traffic discarded or handled on a best-effort basis

Congestion Control in Packet Switched Networks

- send control packet to some or all source nodes
  - requires additional traffic during congestion
- rely on routing information
  - may react too quickly
- end to end probe packets
  - adds to overhead
- add congestion info to packets in transit
  - either backwards or forwards

ATM Traffic Management

- high speed, small cell size, limited overhead bits
- still evolving

reasons existing tools are inadequate for ATM

- majority of traffic not amenable to flow control
- feedback is slow due to reduced transmission time compared with propagation delay
- wide range of application demands
- different traffic patterns
- different network services
- high speed switching and transmission increases volatility

Latency/Speed Effects

- **consider ATM at 150Mbps**
- **takes ~2.8×10^{-6} seconds to insert a single cell**
- **time to traverse network depends on propagation delay and switching delay**
- **assume propagation at two-thirds speed of light**
- **if source and destination on opposite sides of USA, propagation time ~ 48x10-3 seconds**
- **given implicit congestion control, by the time dropped cell notification has reached source, 7.2x106 bits have been transmitted**

Cell Delay Variation

- for ATM voice/video, data is a stream of cells
- delay across network must be short
- rate of delivery must be constant
- there will always be some variation in transit
- delay cell delivery to application so that constant bit rate can be maintained to application

Timing of CBR Cells

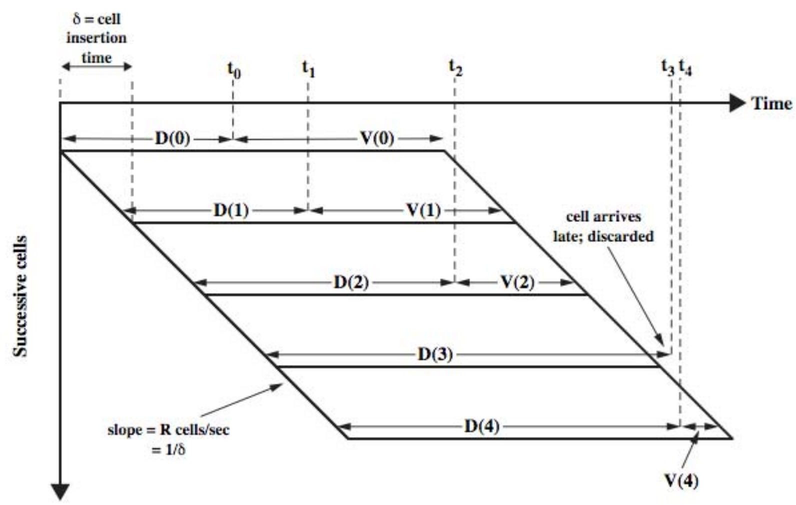

Network Contribution to Cell Delay Variation

- in packet switched networks is due to queuing delays and routing decision time
- in Frame relay networks is similar
- in ATM networks
  - less than frame relay
  - ATM protocol designed to minimize processing overheads at switches
  - ATM switches have very high throughput
  - only noticeable delay is from congestion
  - must not accept load that causes congestion

Cell Delay Variation At The UNI

- application produces data at fixed rate
- 3 layers of ATM processing causes delay
  - interleaving cells from different connections
  - operation and maintenance cell interleaving
  - if using synchronous digital hierarchy frames, these are inserted at physical layer
- cannot predict these delays

Origins of Cell Delay Variation

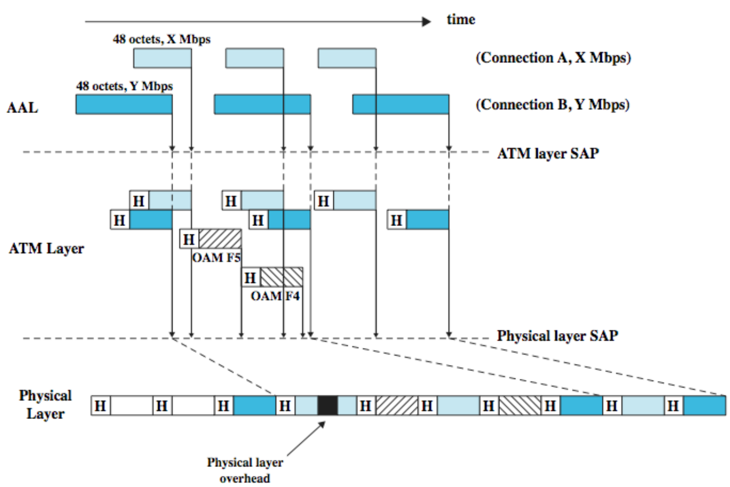

Traffic and Congestion Control Framework

- ATM layer traffic and congestion control should support QoS classes for all foreseeable network services
- should not rely on AAL protocols that are network specific, nor higher level application specific protocols
- should minimize network and end to end system complexity

Traffic Management and Congestion Control Techniques

ATM traffic techniques:

- resource management using virtual paths
- connection admission control
- usage parameter control
- selective cell discard
- traffic shaping

Resource Management Using Virtual Paths

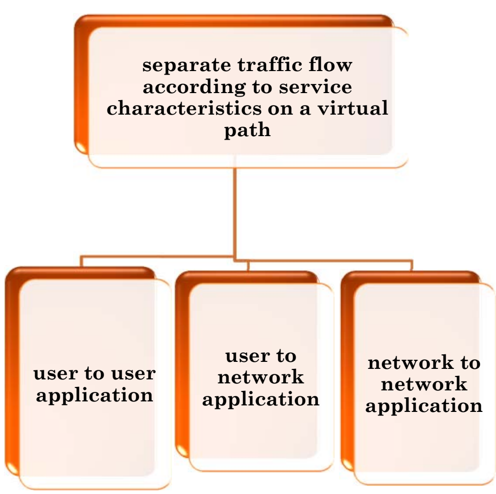

- QoS parameters concerned with are:
  - cell loss ratio
  - cell transfer delay
  - cell delay variation

Timings Considered

- timing intervals considered:
  - cell insertion time
  - round trip propagation time
  - connection duration
  - long term
- traffic control strategy then must:
  - determine whether a given new connection can be accommodated
  - agree performance parameters with subscriber
- now review various control techniques

Resource Management Using Virtual Paths

- separate traffic flow according to service characteristics on a virtual path
  - user to user application
  - user to network application
  - network to network application
- QoS parameters concerned with are:
  - cell loss ratio
  - cell transfer delay
  - cell delay variation

Configuration of VCCs and VPCs

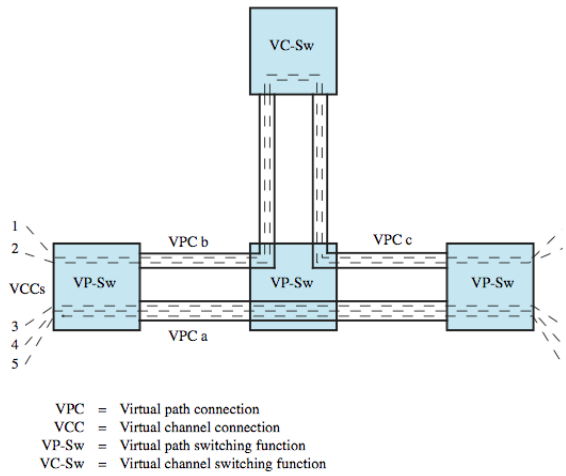

Allocating VCCs within VPC

- all VCCs within VPC should experience similar network performance
- options for allocation:

aggregate peak demand
- set VPC capacity to total of all peak VCC rates
- will meet peak demands, but often underutilized

statistical multiplexing
- set VPC capacity to more than average VCC rates
- will see greater variation but better utilization

Connection Admission Control

- first line of defense
- user specifies traffic characteristics for new connection (VCC or VPC) by selecting a QoS
- network accepts connection only if it can meet the demand and maintain QoS

traffic contract parameters

- ***peak cell rate (PCR)***
- ***cell delay variation (CDV)***
- ***sustainable cell rate (SCR)***
- ***burst tolerance***

Traffic Rate Management

- must discard frames to cope with congestion
  - arbitrarily, no regard for source
  - no reward for restraint so end systems transmit as fast as possible
  - committed information rate (CIR)
    - data in excess of this liable to discard
    - not guaranteed in extreme congestion situations
    - aggregate CIR should not exceed physical data rate
- committed burst size
- excess burst size

Usage Parameter Control (UPC)

- UPC function monitors a connection to ensure traffic obeys contract
- purpose is to protect network resources from overload by one connection
- done on VCC and VPC
- peak cell rate and cell delay variation
- sustainable cell rate and burst tolerance
- UPC discards cells outside traffic contract

Cell Rate Algorithm

- traffic flow is compliant if the peak rate of cell transmission does not exceed the agreed peak cell rate
- cell rate algorithm:

known as the *leaky bucket algorithm*	compliant cells are passed on
defines traffic compliance	noncompliant cells are discarded

Selective Call Discard

- when network at point beyond UPC discards (CLP=1) cells
- aim to discard lower-priority cells when congested to protect higher-priority cells
  - note. can't distinguish between cells originally labeled lower priority, verses those tagged by UPC function

Traffic Shaping

- UPC provides a form of traffic policing
- can be desirable to also shape traffic
- smoothing out traffic flow
- reducing cell clumping
- token bucket

Token Bucket for Traffic Shaping

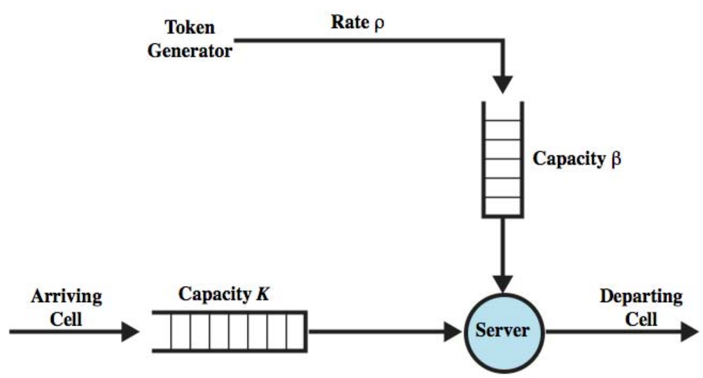

GFR Traffic Management

- guaranteed frame rate (GFR) as simple as UBR from end system viewpoint
- places modest requirements on ATM network
- end system does no policing or shaping of traffic
- may transmit at line rate of ATM adaptor
- no guarantee of delivery
  - so higher layer (eg. TCP) must do congestion control
- user can reserve capacity for each VC
  - ensures application can send at min rate with no loss
  - if no congestion, higher rates maybe used

Frame Recognition

- GFR recognizes frames as well as cells
- when congested, network discards whole frame rather than individual cells
- all cells of a frame have same CLP bit setting
- CLP=1 AAL5 frames lower priority (best effort)
- CLP=0 frames minimum guaranteed capacity

GFR Contract Parameters

- Peak cell rate (PCR)
- Minimum cell rate (MCR)
- Maximum burst size (MBS)
- Maximum frame size (MFS)
- Cell delay variation tolerance (CDVT)

Components of GFR System Supporting Rate Guarantees

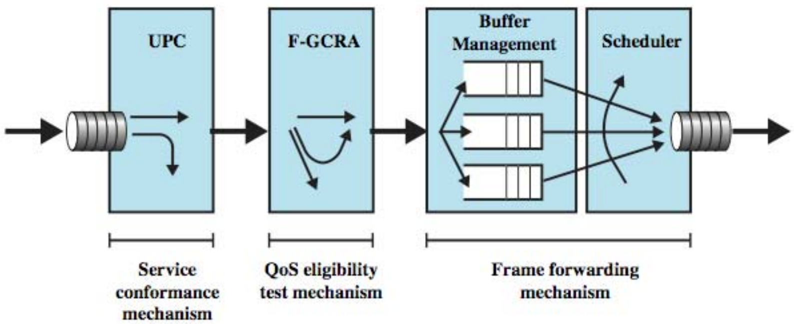

Tagging and Policing

- discriminates between frames that conform to contract and those that don't
- set CLP=1 on all cells in frame if not
  - gives lower priority
- maybe done by network or source
- network may discard CLP=1 cells
  - policing

Buffer Management

- deals with treatment of buffered cells
- congestion indicated by high buffer occupancy
- will discard tagged cells in preference to untagged cells
  - including ones already in buffer to make room
- may do per VC buffering for fairness
- cell discard based on queue-specific thresholds

Scheduling

- preferential treatment to untagged cells
- separate queues for each VC
- make per-VC scheduling decisions
- enables control of outgoing rate of VCs
- VCs get fair capacity allocation
- still meet contract

GFC Conformance Definition

- UPC function monitors each active VC to ensure traffic conforms to contract and tags or discards nonconforming cells
- frame conforms if all cells conform

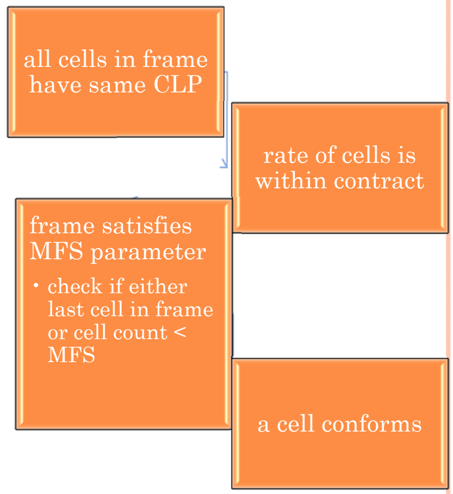

GFR VC Frame Categories

nonconforming frame
- cells of this frame will be tagged or discarded

conforming but ineligible frames
- cells will receive a best-effort service

conforming and eligible frames
- cells will receive a guarantee of delivery

QoS Eligibility Test

- two stage filtering process
- a frame is tested for conformance to contract
  - if not, may discard or tag
  - set upper bound & penalize cells above upper bound
  - do expect attempt to deliver tagged cells
- determine frames eligible for QoS guarantees
  - under GFR contract for VC
  - set lower bound on traffic
  - frames in traffic flow below threshold are eligible

GFR VC Frame Categories

- nonconforming frame
  - cells of this frame will be tagged or discarded
- conforming but ineligible frames
  - cells will receive a best-effort service
- conforming and eligible frames
  - cells will receive a guarantee of delivery
- form of cell rate algorithm is used

SUMMARY

- congestion effects
  - ideal and practical performance
- congestion control
  - backpressure, choke packet, implicit/explicit
- traffic management
  - fairness, QoS, reservations
- ATM traffic management
- ATM-GFR traffic management
  - tagging, policing, buffer, scheduling
- QoS eligibility testing