

NIS/CpE 691CE

Information System Security

Class 9 – 11/18/02

Complex Vulnerabilities in ATM Networks

- Previous discussion focused on identifying weaknesses of a network
- This week, focus on definition of network itself
- IP network example:
 - Robustness is due to packet forwarding - route is not known a priori
 - It is assumed that router forwards packets in general direction of destination using least cost (congestion) routing
 - Intentional discarding of traffic by compromised node violates basic premise of IP network, leading to vulnerabilities
- IP-like networks have been in use for ~30 years - long history has taught the vulnerabilities
- ATM and other networks are relatively new - their underlying vulnerabilities are not yet known.

QoS = f(Source Traffic Statistics)


- Survivability was the original key attraction for IP networks
- ATM networks' guaranteed QoS parameters have been their draw:
 - end-to-end delay
 - jitter
 - cell loss probability
- These parameters are important for voice, video, and other delay sensitive services
- All are guaranteed on a per-call basis if cell-level traffic controls are enforced at UNI
- QoS is based on peak cell rate (PCR) and sustainable cell rate (SCR)
- However:
 - PCR, SCR are hard to obtain - the user parameter is based on PCR, SCR and maximum burstiness of traffic
 - What interval is PCR measured over?
 - The QoS guarantees are not **analytic** guarantees
- A large number of low-bandwidth calls are easier to predict and control.
- Few high-bandwidth calls can strain network severely
 - This is the Central Limit Theorem working against the network operator

Overview of Traffic Controls

- Traffic controls can only be enforced at the edges of the network:
 - The network switches are moving large volumes of traffic at high speeds; measuring and enforcing traffic flows would slow switches down
 - ATM cells are statistically multiplexed on high-capacity links in the core of the network; enforcing per-user controls at this level is problematic
 - The potentially enormous number of QoS contracts between end users and the network cannot be known at every point in the network
- Traffic controls can **only** realistically be enforced at the edge of the network

QoS Underlying Assumptions

- The primary ATM applications are delay-sensitive: cell delay is the primary QoS metric
- Delay depends on traffic volume, buffer architecture and buffer management




Limitations in Measuring PCR

- Peak Cell Rate:

$$PCR = \lim_{T \rightarrow 0} \frac{N}{T}$$

- The minimum interval, T , depends on the time resolution of the system doing the measuring

- Consider:



- $100 \text{ cells} * 53 \text{ bytes} * 8 \text{ bits} / 1 \text{ ms} = 42.4 \text{ Mb/s}$
- $25 \text{ cells} * 53 \text{ bytes} * 8 \text{ bits} / 2.5 \mu\text{s} = 4.24 \text{ Gb/s}$
- Over the shorter measurement window, the PCR is a factor of 100 greater: The user equipment could readily think that it is conforming to its traffic agreement while the network equipment sees a much higher peak rate.


QoS Guarantees and Traffic

- Controlling a traffic source is a key requirement to being able to offer QoS guarantees
- Traffic controls that have been used:
 - resource reservation
 - switch/multiplexor scheduling
 - proactive congestion control
 - call admission control
 - burst admission control
 - peak-rate, sustainable rate policing
 - traffic shaping
 - reactive congestion control
 - window based controls
 - source rate controls
 - assigning cell-loss priorities
 - deliberate blocking of bursts
 - congestion notification


Traffic Characterization

- ATM traffic types:
 - Continuous Bit Rate (CBR) - characterized by a constant rate over long period (e.g., 64 kbps PCM)
 - Variable Bit Rate (VBR) - characterized by its peak bit rate with a lower average bit rate (e.g., compressed video)
 - Unspecified Bit Rate (UBR)

- Modeling burst characteristics:
 - Markov model




- Poisson model



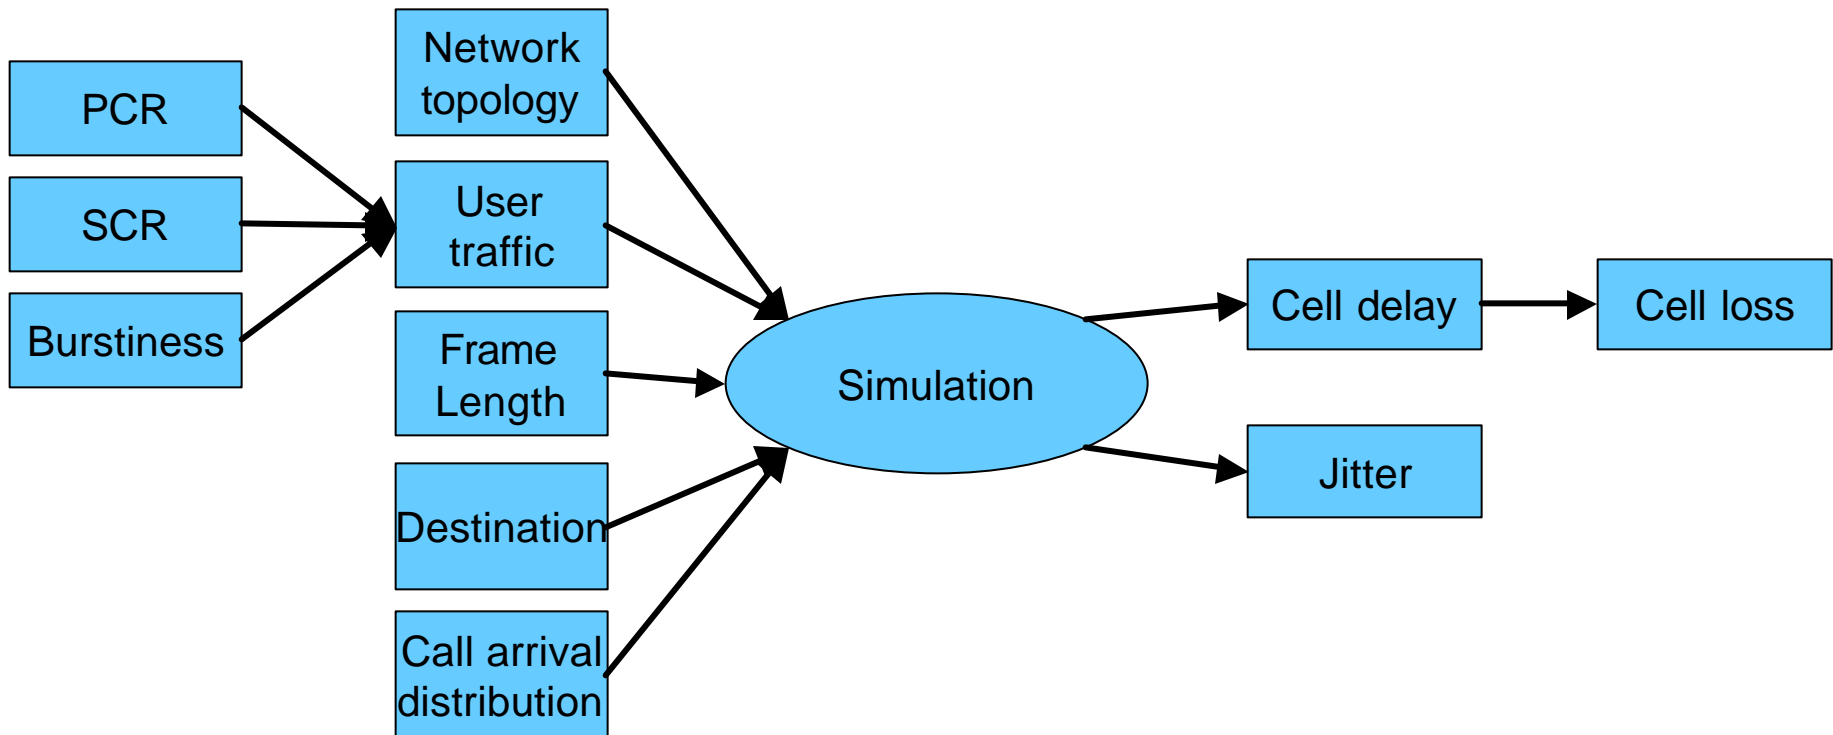
- Limitations: Lack of experimental validation; assumption that source can specify PCR/SCR/burst length; cell level models vs. source level traffic

QoS vs. Traffic Bandwidth Distribution

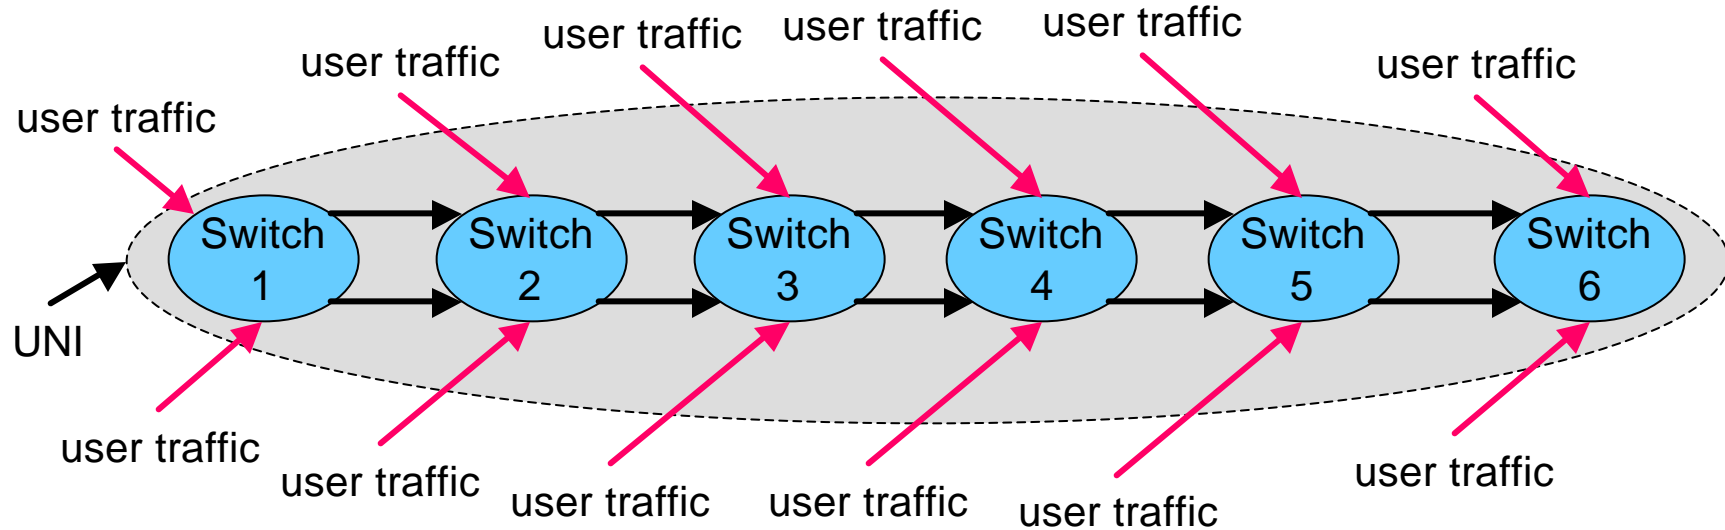
- Network performance:
 - Users are only interested in QoS for **their** traffic
 - Network operator is concerned about
 - link utilization - to recover investment
 - Every users' QoS - for customer retention
- Cell-level traffic controls do not guarantee QoS performance
 - Must look at call-level traffic parameters
 - Focus on traffic bandwidth requirement



ATM Traffic/Utilization Simulation




Network Topology and User Traffic




- User Traffic Model:
 - voice (64 kb/s), video (95 kb/s - 55 Mb/s), and data (50 kb/s - 30 Mb/s)
 - 25 - 2040 traffic sources
 - 80% average link utilization
 - traffic patterns from constant bit rate to highly bursty
 - on/off burstiness
 - variable rate burstiness
 - various frame patterns:
 - variable bit rate with random frame duration
 - variable bit rate with fixed frame duration
 - fixed bit rate with constant inter frame gaps

Traffic Source Models


- Speech - peak bit rate 64 kb/s, average bit rate 25 kb/s



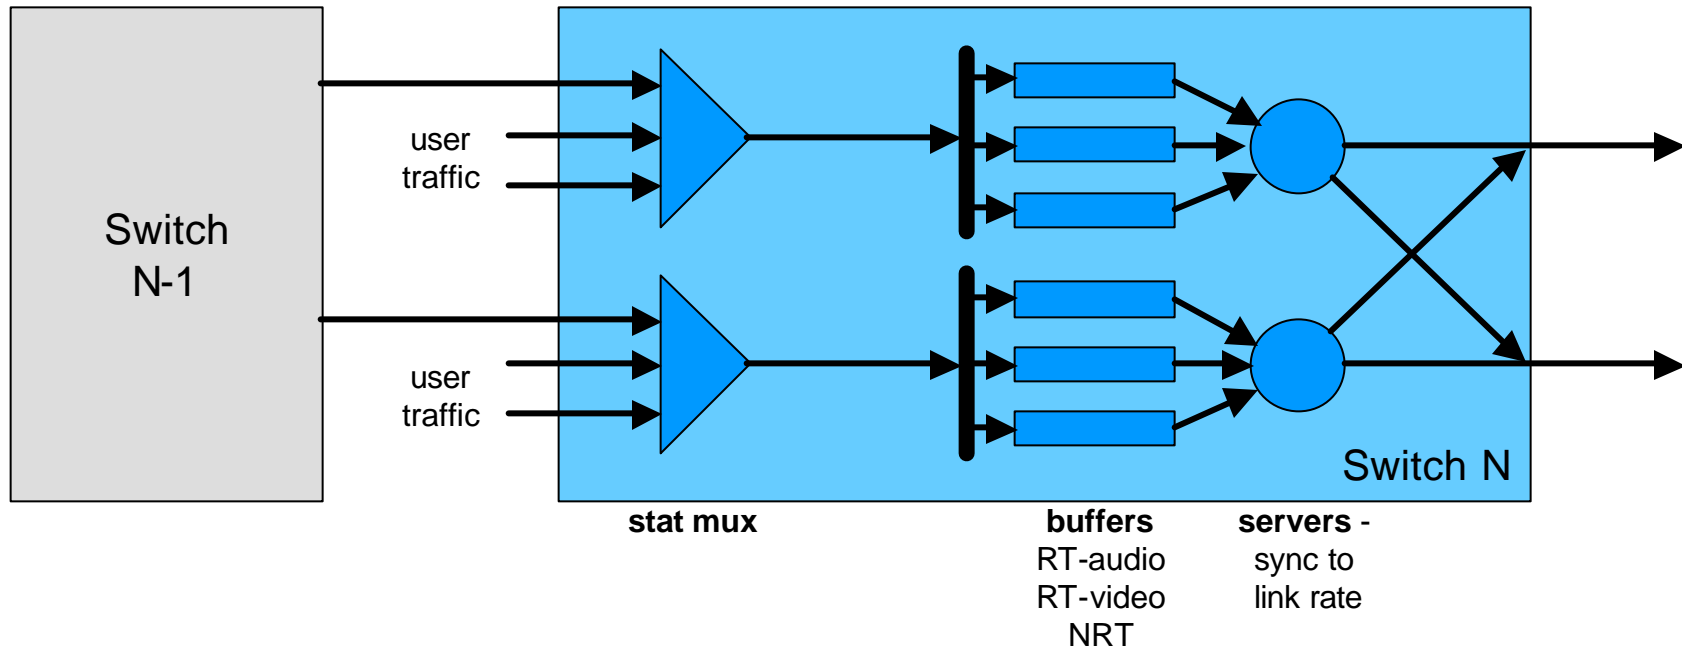
- H.261 video



- Data - peak bit rate 2 Mb/s, average bit rate 1 Mb/s




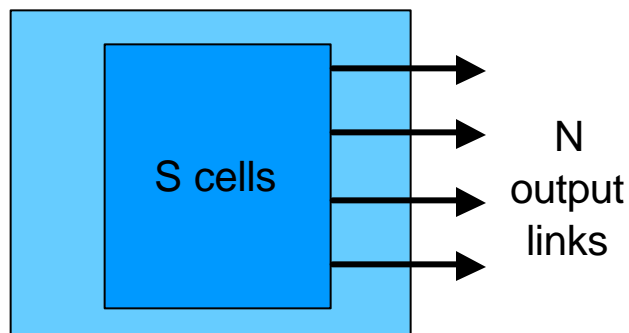
ATM Switch Model



- Arriving cells are assigned to a buffer, depending on their traffic type on space-available basis
- Server
 - gives first priority to real-time traffic over non-real-time,
 - attempts to balance delay between the real-time buffers,
 - periodically services non-real-time traffic

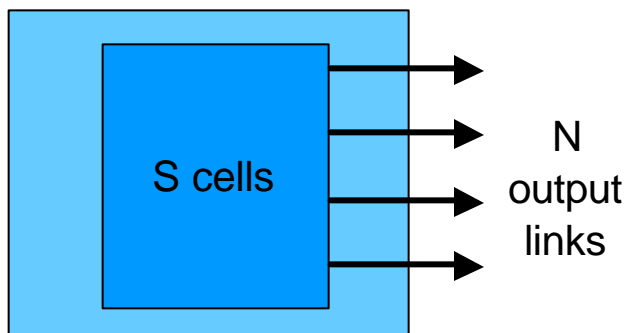
Attacking the ATM Output Buffer

- Assume:
 - To accept a call, adequate resources (e.g., bandwidth) must be available at each link of call
 - User traffic must comply with negotiated parameters at the UNI
- Examine how QoS guarantees can be broken, including severe cell loss
- Underlying issue is due to:
 - Bursty traffic
 - Inability to estimate network traffic
 - Means to determine buffer sizes from traffic
 - Playing on weakness of buffer architecture
- Buffer architectures - fixed at time of switch fabric design:




Attacking the ATM Output Buffer

- Assume:
 - To accept a call, adequate resources (e.g., bandwidth) must be available at each link of call
 - User traffic must comply with negotiated parameters at the UNI
- Examine how QoS guarantees can be broken, including severe cell loss
- Underlying issue is due to:
 - Bursty traffic
 - Inability to estimate network traffic
 - Means to determine buffer sizes from traffic
 - Playing on weakness of buffer architecture
- Buffer architectures - fixed at time of switch fabric design:




**All links are susceptible
to abuse of one**



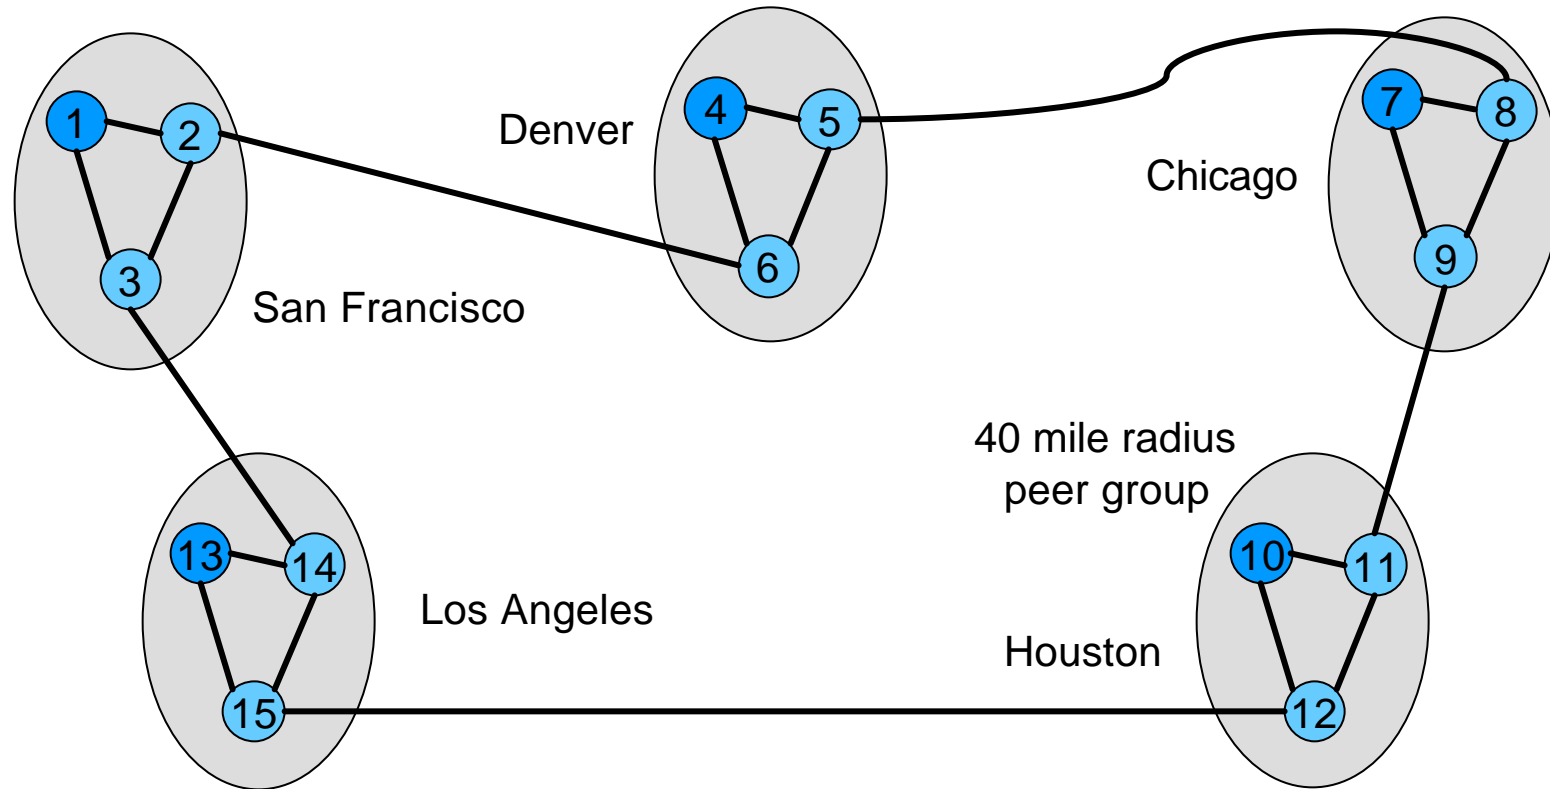
**Susceptible to the
server-splitting problem**

The Buffering Requirement



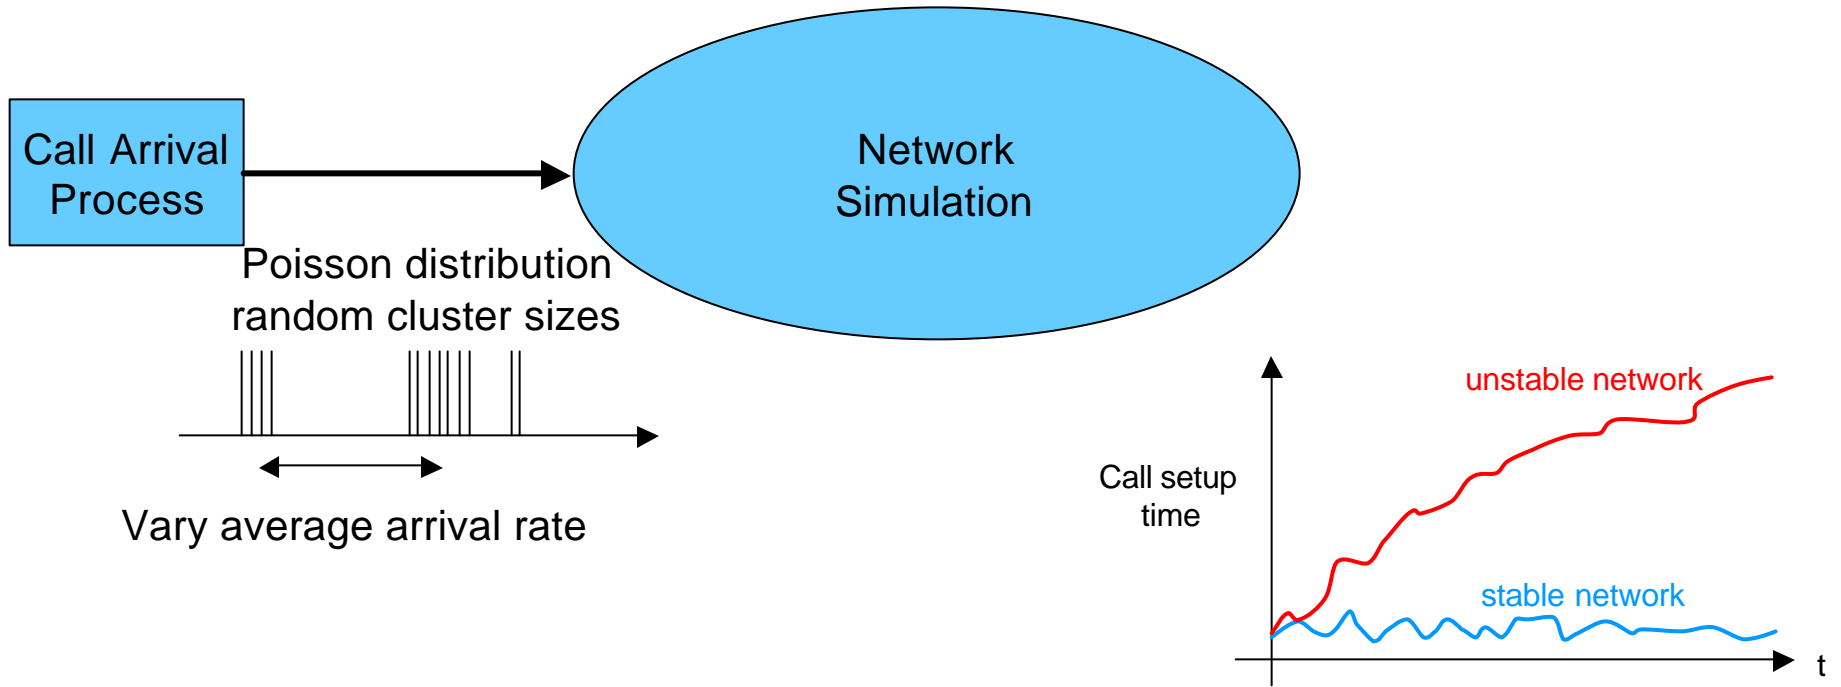
- Example:
 - 8x8 switch with 8000 cell buffer, average burst length = 5, 10% cell loss probability:
 - Single shared buffer sustains 88% of the load
 - Separate buffers sustain 45% of load
 - When burst length (relative to buffer cells per output link) increases, shared buffer may be 80-90% occupied by one output link's traffic - this causes congestion for all links
 - But separate buffers would encounter unacceptable cell loss rates on individual links much earlier

Very High Performance Backbone Network Service (vBNS) - Network Topology



- 15 Node nation-wide network, funded by NSF, implemented by MCI
- N outbound links per node, separate buffers
- To utilize 155.5 Mb/s per link, switch must cell transfer time $2.73\mu\text{s}/N$
- Representative switches provide ~ 13000 cells/buffer

Network Stability Criterion



- Stress network to edge of stability by increasing average call arrival rate

Conclusions

- With acceptable traffic levels at the UNI, internal traffic levels can grow to unacceptable levels
- Cell loss rates up to 40% due to buffer overflow
- Wide variation in throughput across network shows limitations of fixed size buffers
- Simple adjustments in system parameters drive parts of the network to the edge of stability with random traffic patterns.
- Intelligent, concerted effort from geographically dispersed sources to overwhelm a link or node can create even higher cell loss rates
- High bandwidth, bursty traffic can overwhelm buffer capacity without exceeding average link capacity

Homework 8 - due 11/25/02

- Ghosh, p185, exercise 2.