جامعة الأهرام الكندية
Ahram Canadian University

Lecture (10 & 11)
Congestion in Data Networks

Dr. Ahmed ElShafee

Agenda

- Introduction
- Analysis of Congestion effect on Network
- Mechanisms for Congestion Control
- Traffic management

Introduction

Definition

- Congestion occurs when the number of packets being transmitted through a network begins to approach the packet-handling capacity of the network.

The objective of congestion control :

- Keep the number of packets in the network below the level at which performance falls off dramatically. (avoid congesion)

How congestion work?

- A data network or internet is a network of queues.
- At each node (data network switch, internet router), there is a queue of packets for each outgoing channel.
- If the rate at which packets arrive and queue up exceeds the rate at which packets can be transmitted,
- the queue size grows without bound and the delay experienced by a packet goes to infinity.

Example

- when the line -for which packets are queuing -becomes more than 80% utilized,
- the queue length grows at an alarming rate.
- This growth in queue length means increase of node delay.
- Further, since the size of any queue is finite, as queue length grows, eventually the queue must overflow.

Another Example:

- Consider the queuing situation in shown node,
- Any given node has a number of I/O ports attached to it: one or more to other nodes, and zero or more to end systems.
- On each port, packets arrive and depart.

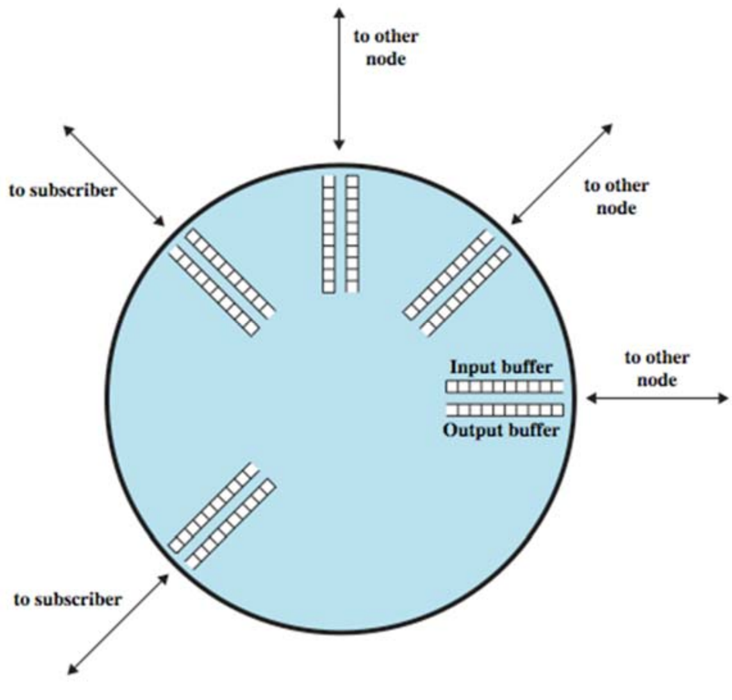

- There are two buffers, or queues, at each port, one to accept arriving packets, and one to hold packets that are waiting to depart. (each has fixed size of variable size as a pool of memory available for all buffering activities and shared in two opposite direction while the buffer sizes is a constant).
- As packets arrive, they are stored in the input buffer of the corresponding port
- The node examines each incoming packet, makes a routing decision, and then moves the packet to the appropriate output buffer as rapidly as possible;

- If packets arrive too fast for the node to process them (make routing decisions) or faster than packets can be cleared from the outgoing buffers, then eventually newly arrived packets will not find any available memory and will be dropped.

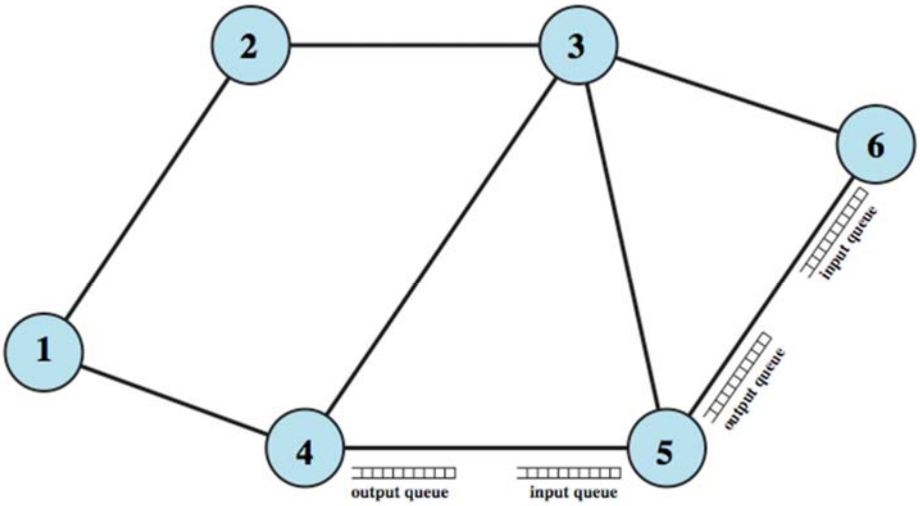

After Congestion scenarios

- When such a saturation point is reached,

1. Node discard any incoming packet for which there is no available buffer space.
2. Node sends flow control over its neighbors so that the traffic flow remains manageable.

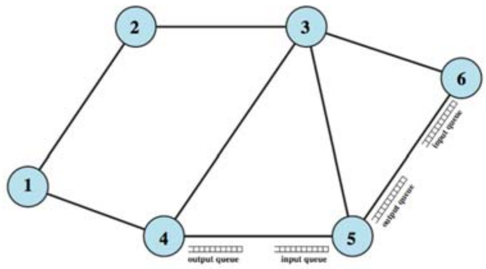

it’s complicated more that it appears!

- each of a node's neighbors is also managing a number of queues.
- If node 6 restrains the flow of packets from node 5, this causes the output buffer in node 5 for the port to node 6 to fill up.
- Thus, congestion at one point in the network can quickly propagate throughout a region or the entire network.

Solution?

- Flow control is indeed in such a way as to manage the traffic on the entire network.

Analysis of Congestion effect on Network

Figures plots normalized load to normalized theoretical throughput

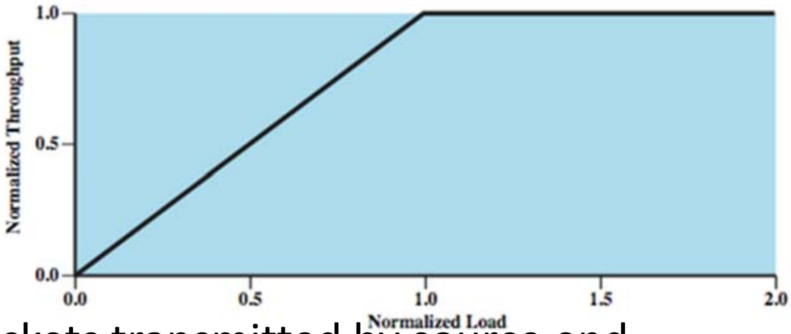

offered load (number of packets transmitted by source end systems),

steady-state total throughput (number of packets delivered to destination end systems)

Both normalized to the maximum theoretical throughput of the network.

For example,

- if a network consists of a single node with two full-duplex 1-Mbps links,
- then the theoretical capacity of the network is 2 Mbps,
- consisting of a 1-Mbps flow in each direction.
- In the ideal case, the throughput of the network increases to accommodate load up to an offered load equal to the full capacity of the network; then normalized throughput remains at 1.0 at higher input loads (ideal performance)

- At normal load, there is some small constant amount of routing and queuing delay.
- As the load on the network increases, queuing delays at each node are added to this fixed amount of delay.
- As each node has a infinite buffer size the network will continue to sustain a normalized throughput of 1.0.
- If packets entering the network is greater than 1.0, while the rate of packets leaving the network is 1.0, queue sizes grow without bound and therefore queuing delays grow without bound.
- That leads to buffer overflow

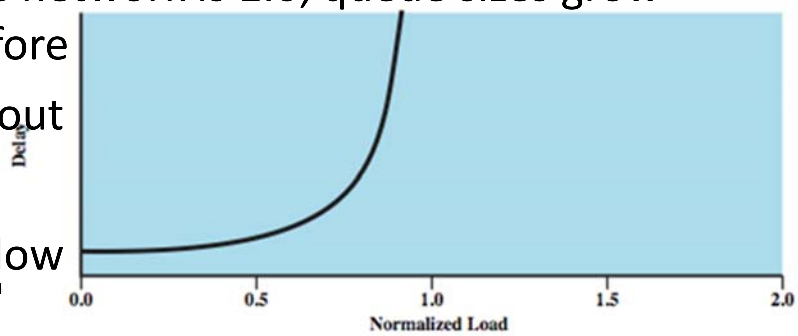

No congestion control

- At light loads, throughput and hence network utilization increases as the offered load increases.

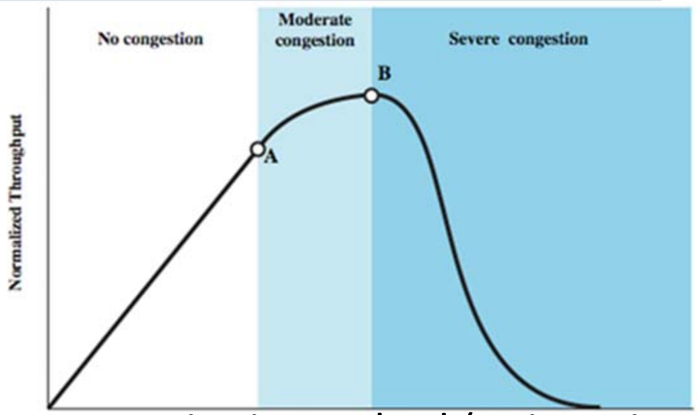

- As the load continues to increase, a point is reached (point A in the plot) beyond which the throughput of the network increases at a rate slower than the rate at which offered load is increased.
- network entry into a moderate congestion state, network continues to cope with the load, although with increased delays.

- As the load on the network continues to increase, the queue lengths of the various nodes continue to grow.
- Eventually, point B is reached where throughput actually drops with increased offered load.

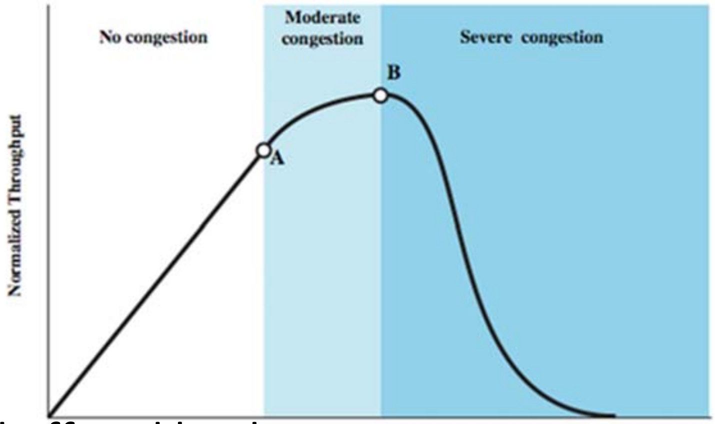

- When the buffers at a node become full, the node must discard packets.
- Thus, the sources must retransmit the discarded packets in addition to new packets.
- Under these circumstances, the effective capacity of the system declines to zero.

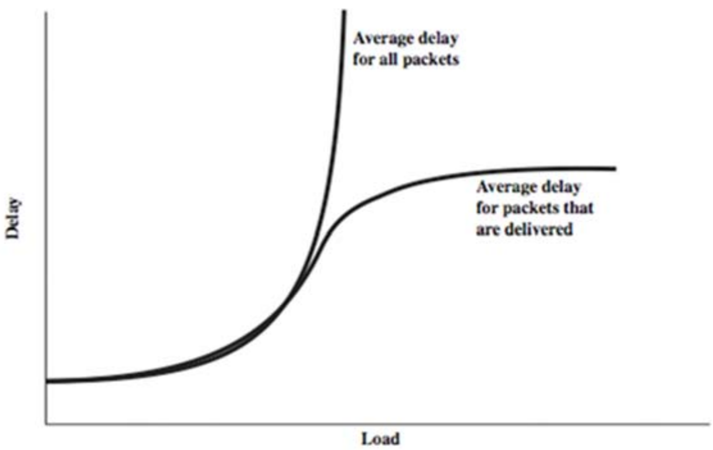

Mechanisms for Congestion Control

We will various techniques for controlling congestion in packet-switching, frame relay, and ATM networks, and in IP-based internets. which include:

- backpressure
- choke packets
- implicit congestion signaling
- explicit congestion signaling

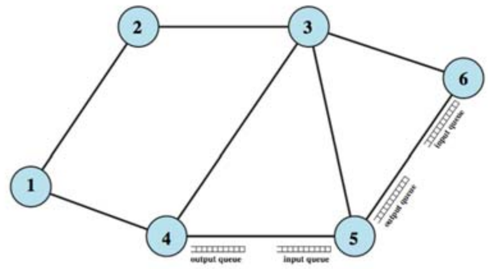

1. Backpressure

Operation:

- if node 6 becomes congested (buffers fill up), then node 6 can slow down or halt the flow of all packets from node 5 (or node 3, or both nodes 5 and 3).
- If this restriction persists, node 5 will need to slow down or halt traffic on its incoming links.
- This flow restriction propagates backward (against the flow of data traffic) to sources, which are restricted in the flow of new packets into the network.

Applied backpressure

- Backpressure can be applied to logical connections, so that the flow from one node to the next is only restricted or halted on some connections, generally the ones with the most traffic.
- the restriction propagates back along the connection to the source.
- It can be used in a connection-oriented network that allows hop-by-hop (from one node to the next) flow control.
- X.25-based packet-switching networks typically provide this feature.

Limitations:

- However, neither frame relay nor ATM has any capability for restricting flow on a hop-by-hop basis.
- In the case of IP-based internets, there have traditionally been no built-in facilities for regulating the flow of data from one router to the next along a path through the internet.

Who's responsibility?

- Each node is responsible of its own congestion

2. Choke Packet

- An example of a choke packet is the ICMP (Internet Control Message Protocol) Source Quench packet.

Strategy 1:

- A choke packet is a control packet generated at a congested node and transmitted back to a source node to restrict traffic flow.
- Either a router or a destination end system may send this message to a source end system, requesting to reduce the sending data rate

- On receipt of a source quench message, the source host should cut back the rate at which it is sending traffic to the specified destination until it no longer receives source quench messages.

Strategy 2:

- In that case, the router or host will issue a source quench message for
  - every datagram that it discards or
  - system may predict congestion and issue source quench messages when its buffers approach capacity.
- In that case, the datagram referred to in the source quench message may well be delivered.
- Receipt of a source quench message does not imply delivery or non-delivery of the corresponding datagram.
- The choke package is a relatively crude technique for controlling congestion.

Who's responsibility?

Network nodes.

Applied chock packet?

connection oriented connection.

3. Implicit Congestion Signaling

The idea

- When network congestion occurs, two things may happen:
  1. The transmission delay for an individual packet from source to destination increases, it is noticeably longer than the fixed propagation delay, and
  2. packets are discarded.
- If a source is able to detect increased delays and packet discards, then it has implicit evidence of network congestion.
- If all sources can detect congestion and, in response, reduce flow on the basis of congestion, then the network congestion will be relieved.

Who's responsibility?

- The responsibility of end systems and does not require action on the part of network nodes.

Applied Implicit Congestion Signaling

- Used in
  - connectionless, or
  - datagram packet-switching networks
  - IP-based internets.
    - In such cases, there are no logical connections through the internet on which flow can be regulated.
  - Connection oriented

Applications

a. TCP

- However, between the two end systems, logical connections can be established at the TCP level.
- TCP uses acknowledging receipt of TCP segments and for regulating the flow of data between source and destination on a TCP connection.
- TCP congestion control techniques based on the ability to detect increased delay and segment loss are discussed in Stallings

b. Frame relay

- For example, in frame relay networks, the LAPF control protocol, which is end to end, includes facilities similar to those of TCP for flow and error control.
- LAPF control is capable of detecting lost frames and adjusting the flow of data accordingly.

4. Explicit Congestion Signaling

General objective

- It is desirable to use as much of the available capacity in a network as possible but still react to congestion in a controlled and fair manner.

Idea

- The network alerts end systems to growing congestion within the network
- the end systems take steps to reduce the offered load to the network.

Who's responsibility?

Network node.

Applied explicit congestion control techniques

- operate over connection-oriented networks and control the flow of packets over individual connections.

Theory of operation

- Explicit congestion signaling approaches can work in one of two directions:

a. **Backward**:

- Notifies the source that congestion avoidance procedures should be initiated where applicable for traffic in the opposite direction of the received notification.
- It indicates that the packets that the user transmits on this logical connection may encounter congested resources.

- **How?** Backward information is transmitted either by
  - altering bits in a header of a data packet headed for the source to be controlled or
  - by transmitting separate control packets to the source.

b. Forward:

- Notifies the user that congestion avoidance procedures should be initiated where applicable for traffic in the same direction as the received notification.
- It indicates that this packet, on this logical connection, has encountered congested resources.
- **How?** this information may be transmitted either as altered bits in data packets or in separate control packets.

So after destination receives congestion signal? What can it do!

- when a forward signal is received by an end system,

1. Send a signal back along the logical connection to the source.
2. the end system is expected to exercise flow control upon the source end system at a higher layer (e.g., TCP).

explicit congestion signaling approaches categories (3):

Binary:

- A bit is set in a data packet as it is forwarded by the congested node.
- When a source receives a binary indication of congestion on a logical connection, it may reduce its traffic flow.

Credit based:

- Destination node provides an explicit credit to a source over a logical connection.
- The credit indicates how many octets or how many packets the source may transmit.
- When the credit is exhausted, the source must await additional credit before sending additional data.
- Credit-based schemes are common for end-to-end flow control, in which a destination system uses credit to prevent the source from overflowing the destination buffers, but credit-based schemes have also been considered for congestion control.

Rate based:

- Destination node provides an explicit data rate limit to the source over a logical connection.
- The source may transmit data at a rate up to the set limit.
- To control congestion, any node along the path of the connection can reduce the data rate limit in a control message to the source.

Traffic management

General idea:

- When a node is saturated and must discard packets, it can apply some simple rule, such as discard the most recent arrival.
- However, other considerations can be used to refine the application of congestion control techniques and discard policy.

Traffic management policies:

1. Fairness

- To ensure that the various flows suffer from congestion equally.

The problem?

- Simply to discard on a last-in-first-discarded basis may not be fair.

The solution?

- For example, a node can maintain a separate queue for each logical connection or for each source-destination pair.
- If all of the queue buffers are of equal length, then queues with the highest traffic load will suffer discards more often, allowing lower-traffic connections a fair share of the capacity.

2. quality of service (QoS)

- During periods of congestion that traffic flows with different requirements be treated differently and provided a different **QoS**.
- For example, a node might transmit higher-priority packets ahead of lower-priority packets in the same queue.
- Or a node might maintain different queues for different QoS levels and give preferential treatment to the higher levels.

- One way to avoid congestion and also to provide assured service to applications is to use a reservation scheme. Such a scheme is an integral part of ATM networks.
- The network agrees to give a defined QoS so long as the traffic flow is within contract parameters;
- excess traffic is either discarded or handled on a best-effort basis, subject to discard.

Summary

- A number of control mechanisms for congestion control in packet-switching networks have been suggested and tried. The following are examples:

1. **Send a control packet from a congested node to some or all source nodes (chock).**

- This choke packet will have the effect of stopping or slowing the rate of transmission from sources and hence limit the total number of packets in the network.
- This approach requires additional traffic on the network during a period of congestion.
-

2. Rely on routing information.

- Routing algorithms, such as ARPANET's, provide link delay information to other nodes, which influences routing decisions.
- This information could also be used to influence the rate at which new packets are produced.
- Because these delays are being influenced by the routing decision, they may vary too rapidly to be used effectively for congestion control.

3. Make use of an end-to-end probe packet.

- Such a packet could be time stamped to measure the delay between two particular endpoints.
- This has the disadvantage of adding overhead to the network.

4. **Allow packet-switching nodes to add congestion information to packets as they go by (implicit and explicit signalling).**

- There are two possible approaches here.
- A node could add such information to packets going in the direction opposite of the congestion.
- This information quickly reaches the source node, which can reduce the flow of packets into the network.
- Alternatively, a node could add such information to packets going in the same direction as the congestion.
- The destination either asks the source to adjust the load or returns the signal back to the source in the packets (or acknowledgments) going in the reverse direction.

The end of course

Wishing you all the best of luck