Congestion Control

- In essence, a packet-switching network is a network of queues.
- At each node, 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.
- Even if the packet arrival rate is less than the packet transmission rate, queue length will grow dramatically as the arrival rate approaches the transmission rate.
- We saw this kind of behavior in Figure 7.16. As a rule of thumb, when the line for which packets are queuing becomes more than 80% utilized, the queue length grows at an alarming rate.

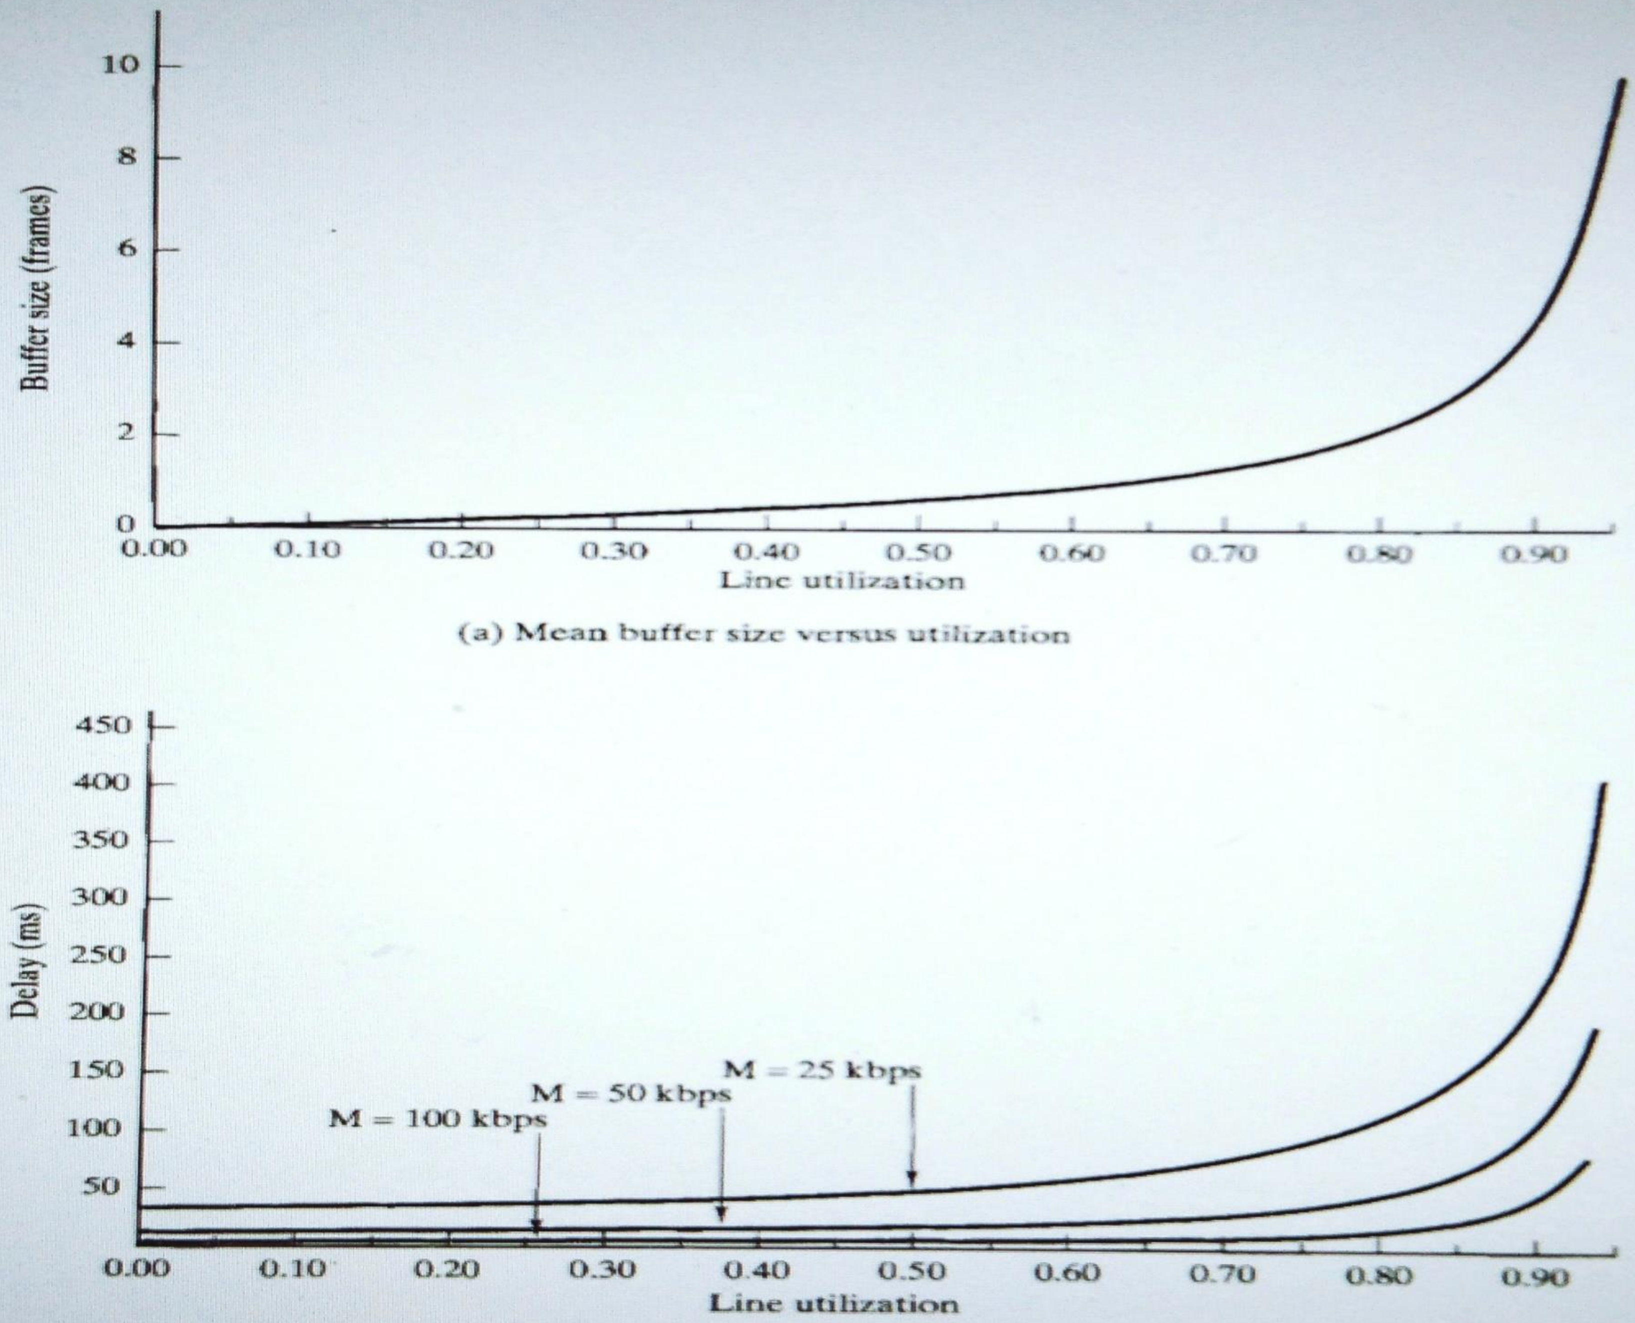

(a) Mean buffer size versus utilization

(b) Mean delay versus utilization

FIGURE 7.16 Buffer size and delay for a statistical multiplexer.

- Consider the queuing situation at a single packet-switching node, such as is illustrated in Figure 9.13.
- Any given node has a number of transmission links attached to it: one or more to other packet-switching nodes, and zero or more to host systems.
- On each link, packets arrive and depart.
- We can consider that there are two buffers at each link, one to accept arriving packets, and one to hold packets that are waiting to depart.

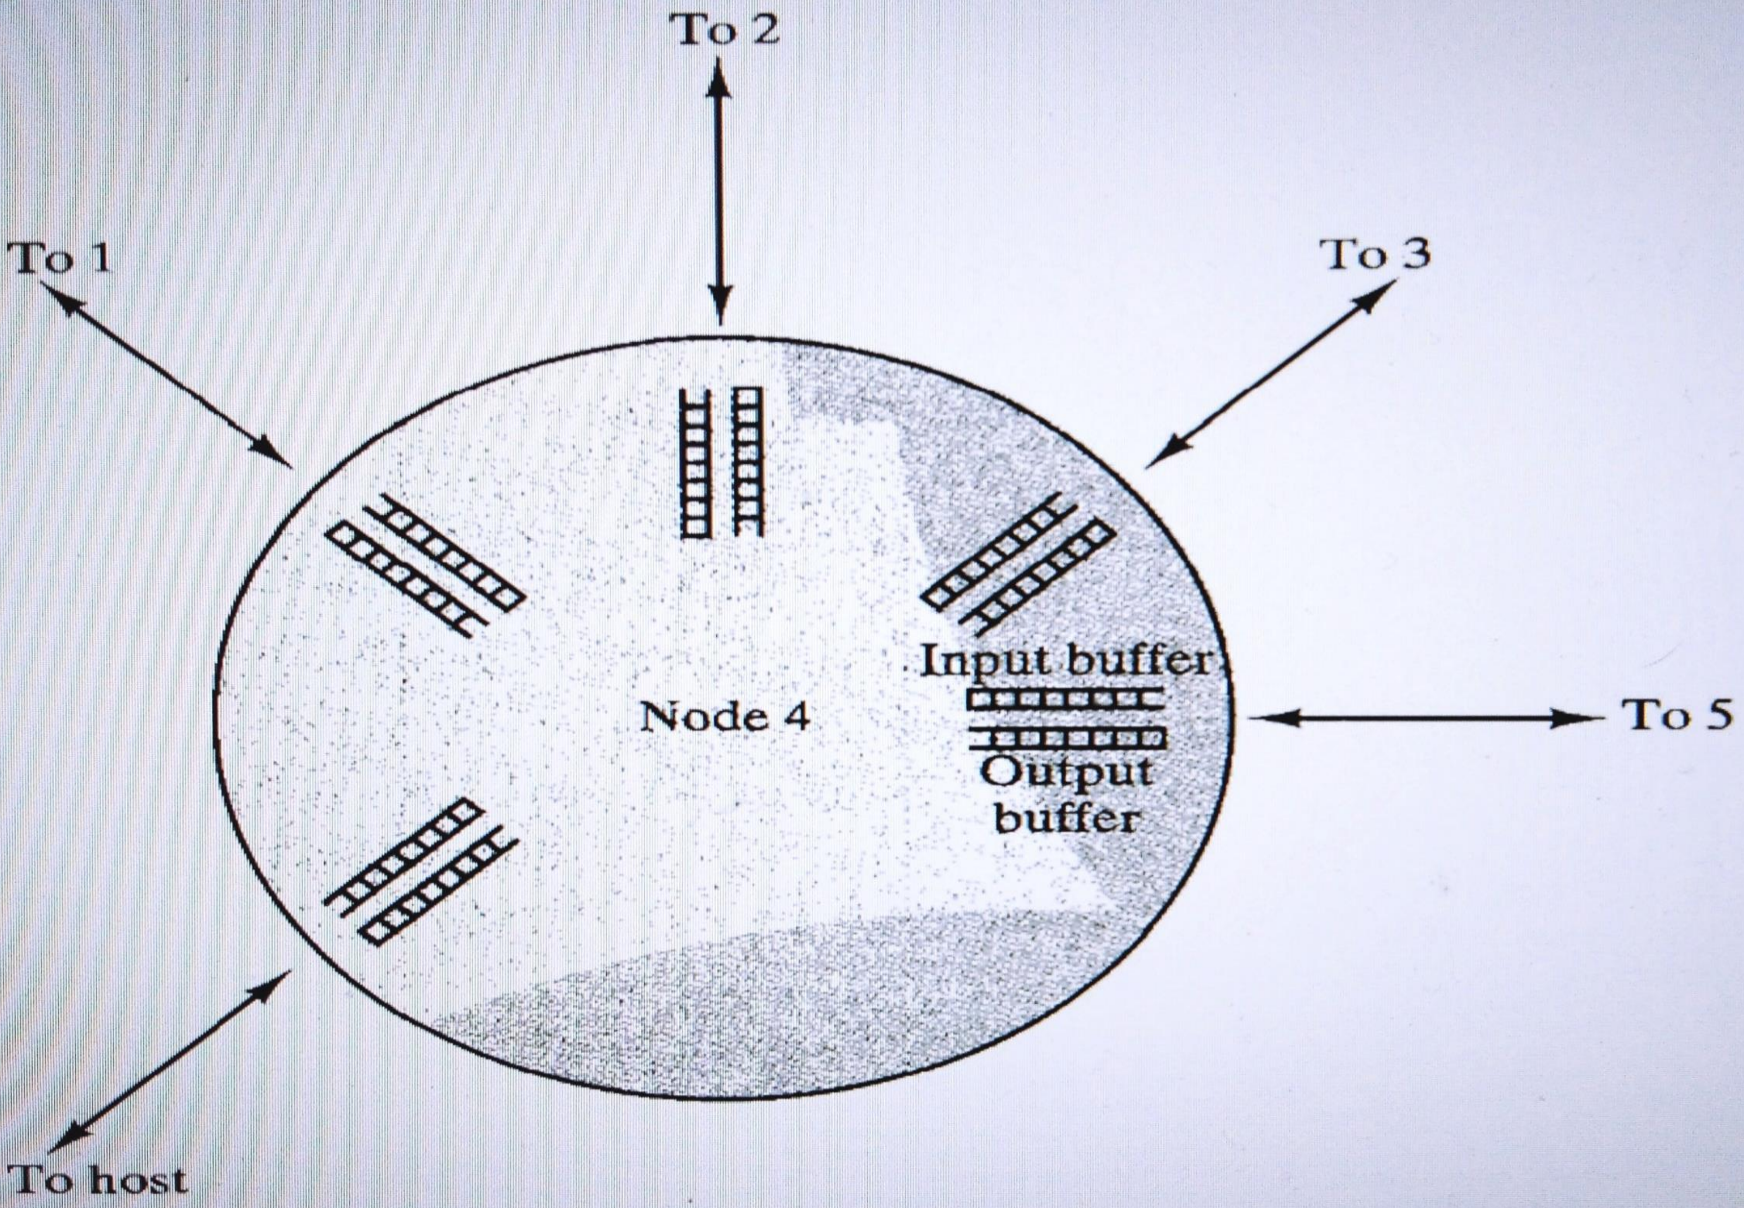

FIGURE 9.13 Input and output queues at node 4 of Figure 9.5.

- In any case, as packets arrive, they are stored in the input buffer of the corresponding link.
- The node examines each incoming packet to make a routing decision, and then moves the packet to the appropriate output buffer.
- Packets queued up for output are transmitted as rapidly as possible; this is, in effect, statistical time division multiplexing.
- Now, 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, packets will arrive for which no memory is available.

- When such a saturation point is reached, one of two general strategies can be adopted.
- The first such strategy is to simply discard any incoming packet for which there is no available buffer space.
- The alternative is for the node that is experiencing these problems to exercise some sort of flow control over its neighbors so that the traffic flow remains manageable.
- But, as Figure 9.14 illustrates, 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 link to node 6 to fill up.
- Thus, congestion at one point in the network can quickly propagate throughout a region or throughout all of the network.
- While flow control is indeed a powerful tool, we need to use it in such a way as to manage the traffic on the entire network.

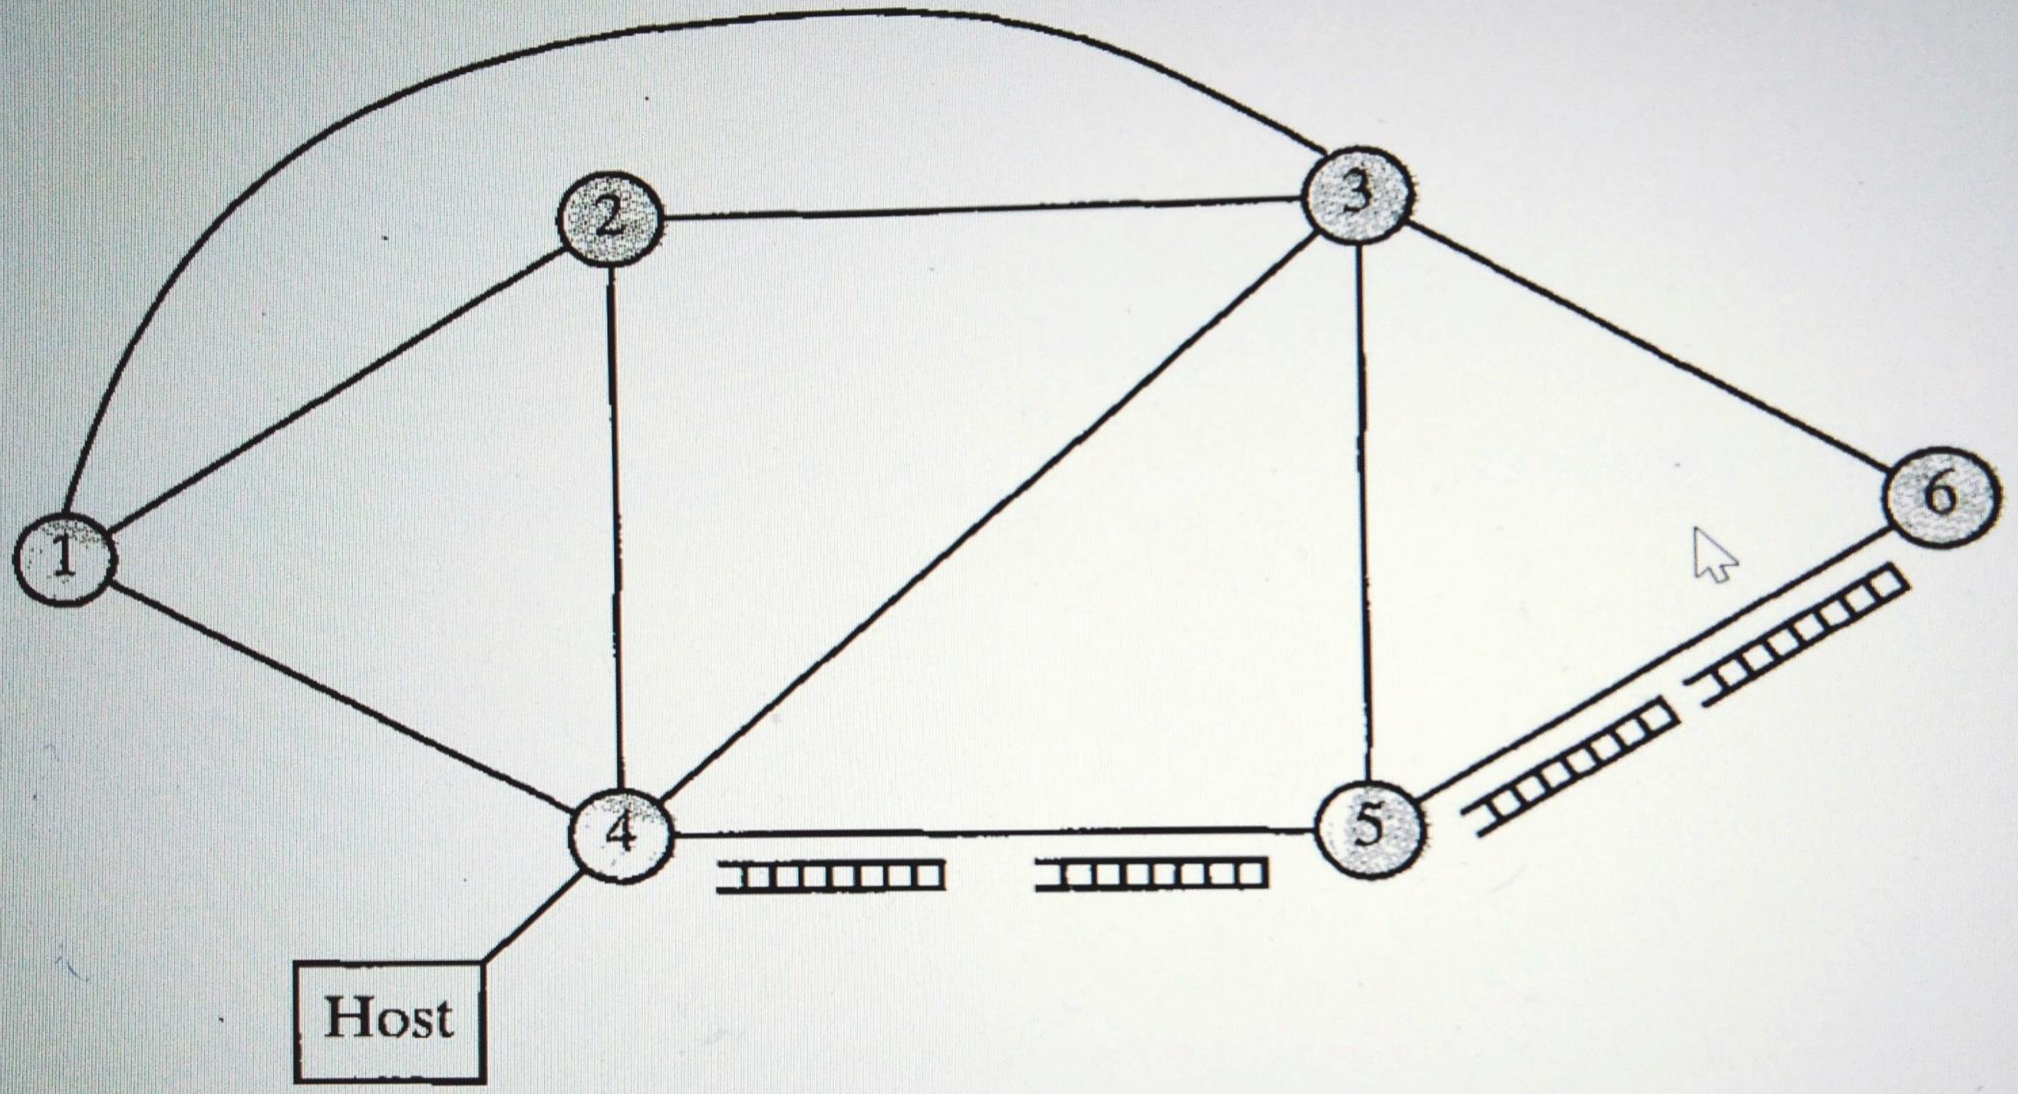

FIGURE 9.14 The interaction of queues in a packet-switching network.

- Figure 9.15 shows the effect of congestion in general terms.
- Figure 9.15a plots the throughput of a network (number of packets delivered to destination stations) versus the offered load (number of packets transmitted by source stations).
- Both axes are normalized to the maximum capacity of the network, which can be expressed as the rate at which the network is theoretically capable of handling packets.
- In the ideal case, throughput and, hence, network utilization increase to accommodate an offered load up to the maximum capacity of the network.
- Utilization then remains at 100%.

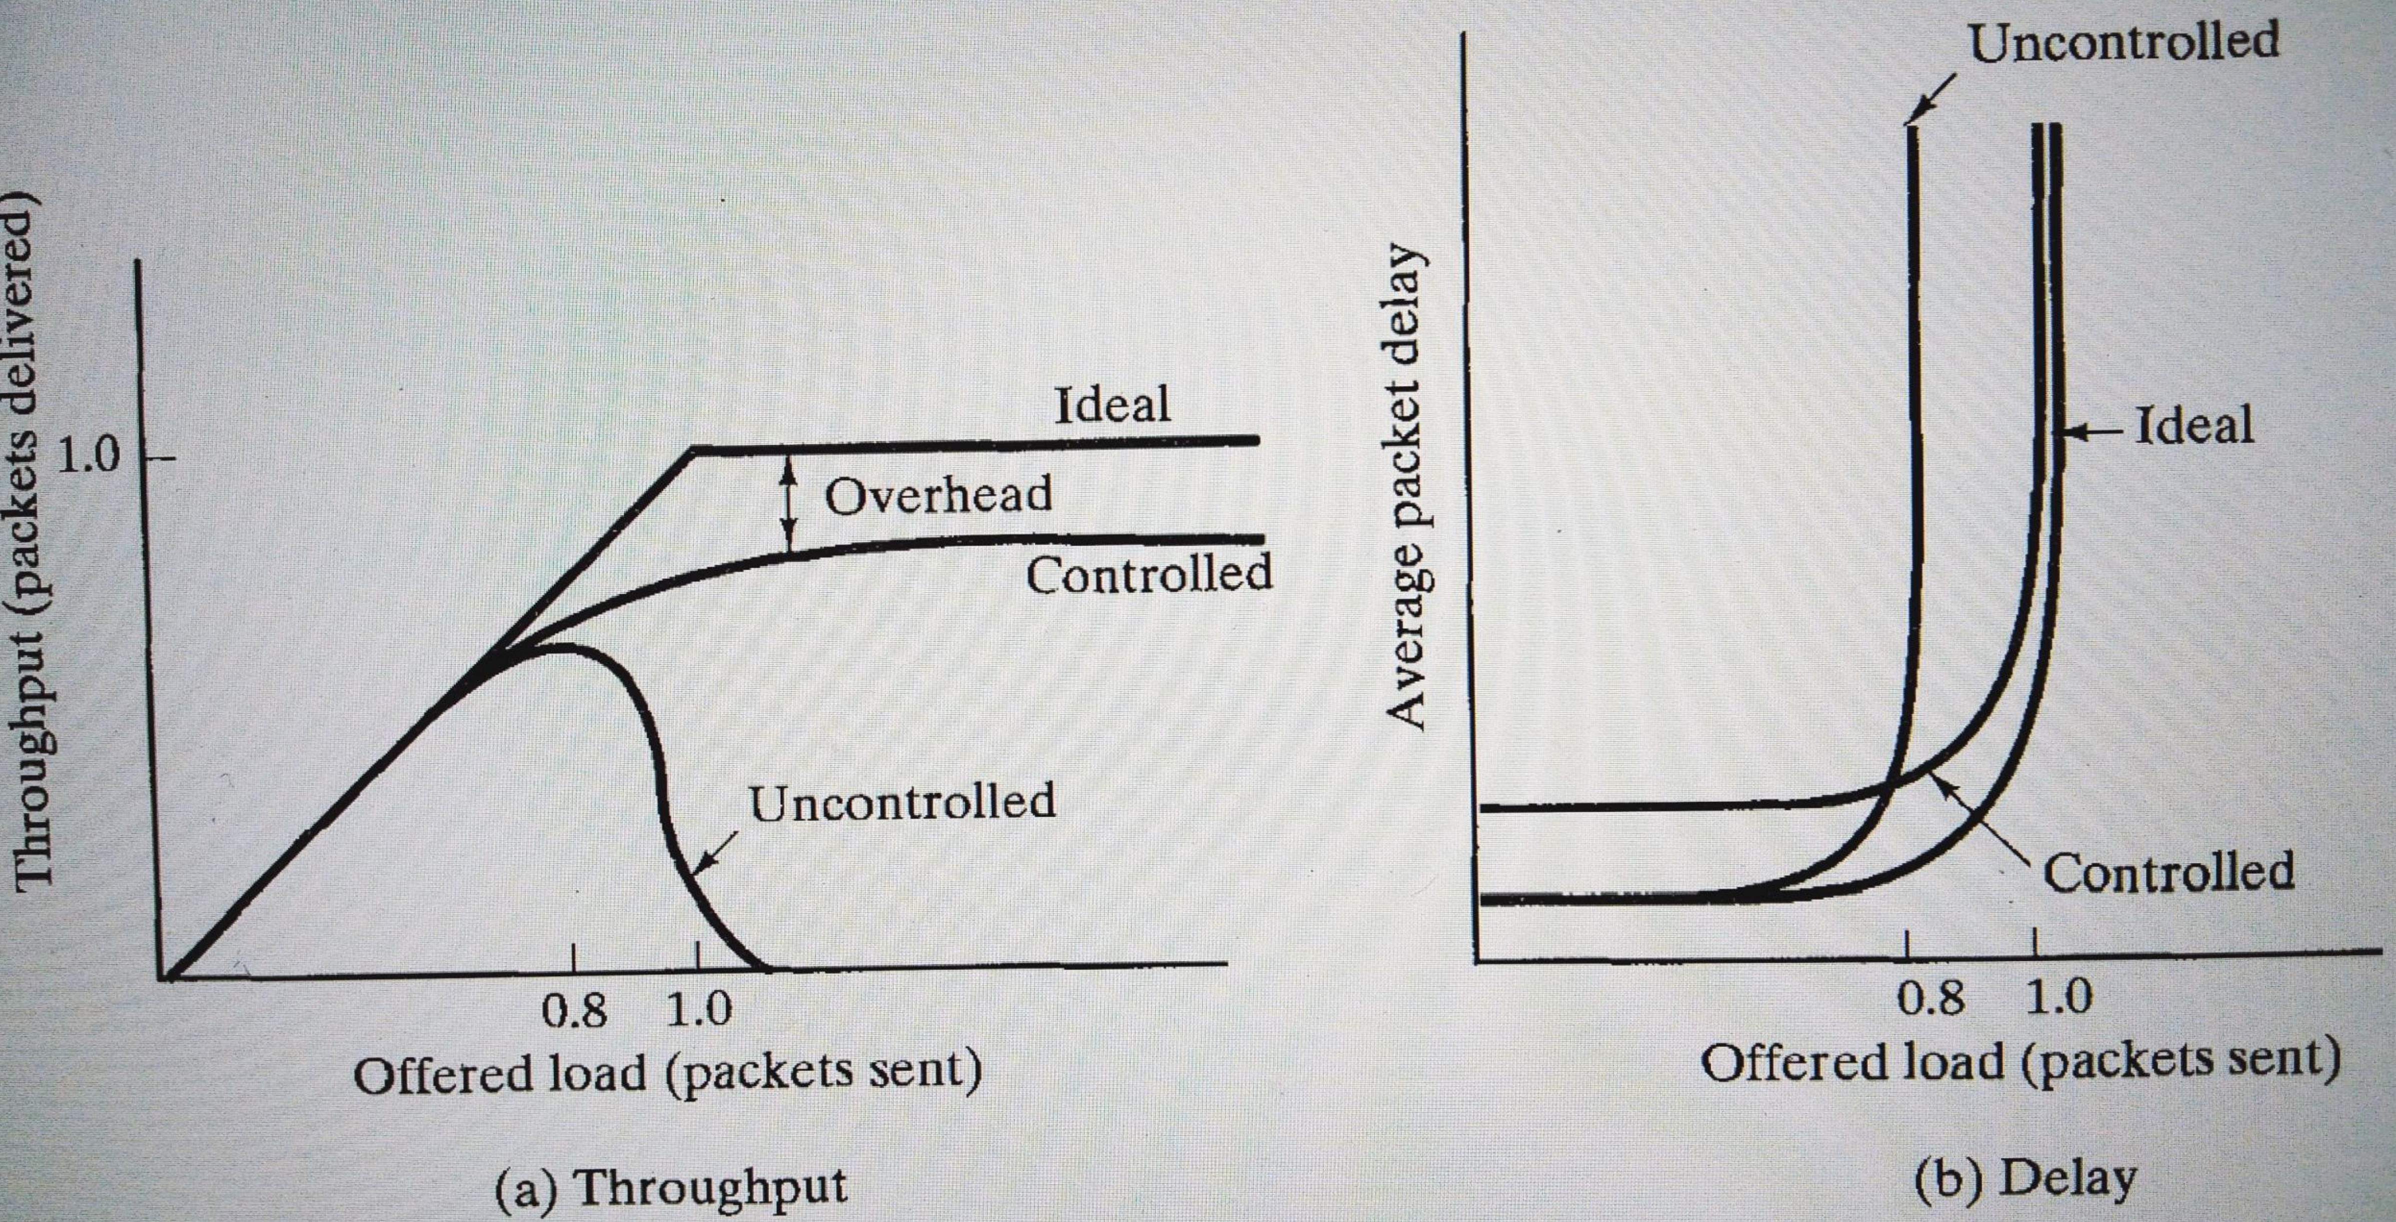

FIGURE 9.15 The effects of congestion.

- The ideal case, of course, requires that all stations somehow know the timing and rate of packets that can be presented to the network, which is impossible.
- If no congestion control is exercised, we have the curve labeled "uncontrolled."
- As the load increases, utilization increases for a while.
- Then as the queue lengths at the various nodes begin to grow, throughput actually drops because the buffers at each node are of finite size.
- When a node's buffers are full, it must discard packets.

- Thus, the source stations must retransmit the discarded packets in addition to the new packets; this only exacerbates the situation:
- As more and more packets are retransmitted, the load on the system grows, and more buffers become saturated.
- While the system is trying desperately to clear the backlog, stations are pumping old and new packets into the system.
- Even successfully delivered packets may be retransmitted because it takes so long to acknowledge them:
- The sender assumes that the packet did not go through.
- Under these circumstances, the effective capacity of the system is virtually zero.

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. 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 ARPANETs, 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 timestamped to measure the delay between two particular endpoints. This procedure has the disadvantage of adding overhead to the network.
4. 4. Allow packet-switching nodes to add congestion information to packets as they go by. 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.