

Communication Networks: Sharing and Switches

So far in this course we have studied techniques to engineer a *point-to-point* communication link. Two important considerations were at the heart of most everything we studied:

1. *Improving link communication reliability*: Inter-symbol interference (ISI) and noise conspired to introduce errors in transmission. We developed techniques to select a suitable sampling rate and method using *eye diagrams* and then reduced the bit-error rate using *channel coding* (block, Reed-Solomon, and convolutional codes).
2. *Sharing a link*: We wanted to share the same communication medium amongst k different receivers, each tuned to a different frequency. To achieve this goal, we used *digital modulation*, and learned that understanding the frequency response of an LTI system and designing filters are key building blocks for this task.

We now turn to the study of communication networks—systems that connect three or more computers (or phones)¹ together.

The key idea that we will use to engineer communication networks is *composition*: we will build small networks by composing links together, and build larger networks by composing smaller networks together.

The fundamental challenges in the design of a communication network are the same as those that face the designer of a communication link: *sharing* and *reliability*. The big difference is that the sharing problem is considerably more challenging, and many more things can go wrong in networking² many computers together, making communication more unreliable than a single link's unreliability. The next few lectures will show you these challenges and you will understand the key ideas in how these challenges are overcome.

In addition to sharing and reliability, an important and difficult problem that many communication networks (such as the Internet) face is *scalability*: how to engineer a very large, global system. We won't say very much about scalability in this course, leaving this important issue to future courses in EECS.




Figure 15-1: A communication network with a link between every pair of computers has a quadratic number of links. This approach is too expensive and is especially untenable when the computers are far from each other.

15.1 Sharing with Switches

The collection of techniques used to design a communication link, including modulation and error-correcting channel coding, is usually implemented in a module called the *physical layer* (or “PHY” for short). The sending PHY takes a stream of bits and arranges to send it across the link to the receiver; the receiving PHY provides its best estimate of the stream of bits sent from the other end. On the face of it, once we know how to develop a communication link, connecting a collection of N computers together is ostensibly quite straightforward: one could simply connect each pair of computers with a wire and use the physical layer running over the wire to communicate between the two computers. This picture for a small 5-node network is shown in Figure 15-1.

This simple strawman using dedicated pairwise links has two severe problems. First, it is extremely expensive. The reason is that the number of distinct communication links that one needs to build scales quadratically with N , because there are about N^2 different bi-directional links in this design. The cost of operating such a network would be prohibitively expensive, and each additional node added to the network would incur a cost proportional to the size of the network! Second, some of these links would have to span an enormous distance; imagine how the computers in Cambridge, MA, would be connected to those in Cambridge, UK, or (to go further) to those in India or China. Such “long-haul” links are difficult to engineer, so one can’t assume that they will be available in abundance. Clearly we need a better design, one that can “do for a dime what any fool can do for a dollar”.³ The key to a practical design of a communication network is a special computing device called a *switch*. A switch has multiple “interfaces” (or ports) on it; a link (wire or radio) can be connected to each interface. The switch allows multiple different communications between different pairs of computers to run over each individual link—that is, it arranges for the network’s links to be *shared* by different communications. In addition to




Figure 15-2: A simple network topology showing communicating end points, links, and switches.

the links, the switches themselves have some resources (memory and computation) that will be shared by all the communicating computers.

Figure 15-2 shows the general idea. A switch receives bits that are encapsulated in *data frames* (in some networks, these are called *packets*, as explained below) arriving over its links, processes them (in a way that we will make precise later), and forwards them (again, in a way that we will make precise later) over one or more other links.

We will use the term *end points* to refer to the communicating computers, and call the switches and links over which they communicate the *network infrastructure*. The resulting structure is termed the *network topology*, and consists of *nodes* (the switches and end points) and links. A simple network topology is shown in Figure 15-2. We will model the network topology as a *graph* to solve various problems.

Figure 15-3 show a few switches of relatively current vintage (ca. 2006).

15.1.1 Three Problems That Switches Solve

The fundamental functions performed by switches are to multiplex and demultiplex data frames belonging to different computer-to-computer information transfer sessions, and to determine the link(s) along which to forward any given data frame. This task is essential because a given physical link will usually be shared by several concurrent sessions between different computers. We break these functions into three problems:

1. **Forwarding:** When a data frame arrives at a switch, the switch needs to process it, determine the correct outgoing link, and decide when to send the frame on that link.
2. **Routing:** Each switch somehow needs to determine the topology of the network, so that it can correctly construct the data structures required for proper forwarding. The process by which the switches in a network collaboratively compute the network topology, adapting to various kinds of failures, is called routing. It does not happen on each data frame, but occurs in the “background”.
3. **Resource allocation:** Switches allocate their resources—access to the link and local memory—to the different communications that are in progress.




Figure 15-3: A few modern switches.

Over time, two radically different methods have been developed for solving these problems. These techniques differ in the way the switches forward data and allocate resources (there are also some differences in routing, but they are less significant). The first method, used by networks like the telephone network, is called *circuit switching*. The second method, used by networks like the Internet, is called *packet switching*.

There are two crucial differences between the two methods, one philosophical and the other mechanistic. The mechanistic difference is the easier one to understand, so we'll talk about it first. In a circuit-switched network, the frames do not (need to) carry any special information that tells the switches how to forward information, while in packet-switched networks, they do. The philosophical difference is more substantive: a circuit-switched network provides the abstraction of a *dedicated link* of some bit rate to the communicating entities, whereas a packet switched network does not.⁴ Of course, this dedicated link traverses multiple physical links and at least one switch, so the end points and switches must do some additional work to provide the illusion of a dedicated link. A packet-switched network, in contrast, provides no such illusion; once again, the end points and switches must do some work to provide reliable and efficient communication service to the applications running on the end points.

15.2 Circuit Switching

The transmission of information in circuit-switched networks usually occurs in three phases (see Figure 15-4):




Figure 15-4: Circuit switching requires setup and teardown phases.




Figure 15-5: Circuit switching with Time Division Multiplexing (TDM). Each color is a different conversation and there are a maximum of $N=6$ concurrent communications on the link in this picture. Each communication (color) is sent in a fixed time-slot, modulo N .

1. the *setup phase*, in which some state is configured at each switch along a path from source to destination,
2. the *data transfer phase* when the communication of interest occurs, and
3. the *teardown phase* that cleans up the state in the switches after the data transfer ends.

Because the frames themselves contain no information about where they should go, the setup phase needs to take care of this task, and also configure (reserve) any resources needed for the communication so that the illusion of a dedicated link is provided. The teardown phase is needed to release any reserved resources.

15.2.1 Example: Time-Division Multiplexing (TDM)

A common (but not the only) way to implement circuit switching is using *time-division multiplexing (TDM)*, also known as *isochronous transmission*. Here, the physical capacity, or

bit rate,⁵ of a link connected to a switch, C (in bits/s), is conceptually broken into some number N of virtual “channels” (or time slots), such that the ratio C/N bits/s is sufficient for each information transfer session (such as a telephone call between two parties). Call this ratio, R , the *rate* of each independent transfer session. Now, if we constrain each frame to be of some fixed size, s bits, then the switch can perform time multiplexing by allocating the link’s capacity in time-slots of length s/C units each, and by associating the i th time- slice to the i th transfer (modulo N). It is easy to see that this approach provides each session with the required rate of R bits/s, because each session gets to send s bits over a time period of Ns/C seconds, and the ratio of the two is equal to $C/N = R$ bits/s.

Each data frame is therefore forwarded by simply using the time slot in which it arrives at the switch to decide which port it should be sent on. Thus, the state set up during the first phase has to associate one of these channels with the corresponding soon-to-follow data transfer by allocating the i th time-slice to the i th transfer. The end points transmitting data send frames only at the specific time-slots that they have been told to do so by the setup phase.

Other ways of doing circuit switching include *wavelength division multiplexing (WDM)*, *frequency division multiplexing (FDM)*, and *code division multiplexing (CDM)*; the latter two (as well as TDM) are used in some wireless networks, while WDM is used in some high- speed optical networks.

15.2.2 Pros and Cons

Circuit switching makes sense for a network where the workload is relatively uniform, with all information transfers using the same capacity, and where each transfer uses a *constant bit rate* (or near-constant bit rate). The most compelling example of such a workload is telephony, where each digitized voice call might operate at 64 kbits/s. Switching was first invented for the telephone network, well before computers were on the scene, so this design choice makes a great deal of sense. The classical telephone network as well as the cellular telephone network in most countries still operate in this way, though telephony over the Internet is becoming increasingly popular and some of the network infrastructure of the classical telephone networks is moving toward packet switching.

However, circuit-switching tends to waste link capacity if the workload has a *variable bit rate*, or if the frames arrive in bursts at a switch. Because a large number of computer applications induce burst data patterns, we should consider a different link sharing strategy for computer networks. Another drawback of circuit switching shows up when the $(N + 1)^{\text{st}}$ communication arrives at a switch whose relevant link already has the maximum number (N) of communications going over it. This communication must be denied

⁵This number is sometimes referred to as the “bandwidth” of the link. Technically, bandwidth is a quantity measured in Hertz and refers to the width of the frequency over which the transmission is being done. To avoid confusion, we will use the term “bit rate” to refer to the number of bits per second that a link is currently operating at, but the reader should realize that the literature often uses “bandwidth” to refer to this term. The reader should also be warned that some people (curmudgeons?) become apoplectic when they hear someone using “bandwidth” for the bit rate of a link. A more reasonable position is to realize that when the context is clear, there’s not much harm in using “bandwidth”. The reader should also realize that in practice most wired links usually operate at a single bit rate (or perhaps pick one from a fixed set when the link is configured), but that wireless links using radio communication can operate at a range of bit rates, adaptively selecting the modulation and coding being used to cope with the time-varying channel conditions caused by interference and movement.

access (or admission) to the system, because there is no capacity left for it. For applications that require a certain minimum bit rate, this approach might make sense, but even in that case a “busy tone” is the result. However, there are many applications that don’t have a minimum bit rate requirement (email and file downloads are examples); for this reason as well, a different sharing strategy is worth considering.

Packet switching doesn’t have these drawbacks.

15.3 Packet Switching

The best way to overcome the above inefficiencies is to allow for any sender to transmit data at any time, but yet allow the link to be shared. Packet switching is a way to accomplish this task, and uses a tantalizingly simple idea: add to each frame of data a little bit of information that tells the switch how to forward it. This information is added in the form of a *packet header* and the resulting frame is called a *packet*.⁶ In the most common form of packet switching, the header of each packet contains the *address* of the destination, which uniquely identifies the destination of data. The switches use this information to process and forward each packet. Packets usually also include the sender’s address to help the receiver send messages back to the sender.

The job of the switch is to use the destination address as a key and perform a lookup on a data structure called a *routing table* (or *forwarding table*; the distinction between the two is sometimes important and will become apparent in a later lecture in the course). This lookup returns an outgoing link to forward the packet on its way toward the intended destination.

While forwarding is a relatively simple lookup in a data structure, the trickier question that we will spend time on is determining how the entries in the routing table are obtained. The plan is to use a background process called a *routing protocol*, which is typically implemented in a distributed manner by the switches. There are several types of routing protocols that one might consider, and we will study two common classes of protocols in later lectures. For now, it is enough to understand that the result of running a routing protocol is to obtain routes (which you can think of as paths for the time being) in the network to every destination—each switch dynamically constructs and updates its routing table using the routing protocol and uses this table to forward data packets.

Switches in packet-switched networks that implement the functions described in this section are often called *routers*. Packet forwarding and routing using the Internet Protocol (IP) in the Internet is an example of packet-switching.

15.3.1 Why Packet Switching Works: Statistical Multiplexing

Packet switching does not provide the illusion of a dedicated link to any pair of communicating end points, but it has a few things going for it:

1. it doesn’t waste the capacity of any link because each switch can send any packet available to it that needs to use that link,
2. it does not require any setup or teardown phases and so can be used even for small transfers without any overhead, and

⁶Sometimes, the term *datagram* is used instead of (or in addition to) the term “packet”.




Figure 15-6: Packet switching works because of statistical multiplexing. This picture shows a simulation of N senders, each connected at a fixed bit rate of 1 megabit/s to a switch, sharing a single outgoing link. The y-axis shows the aggregate bit rate (in megabits/s) as a function of time (in milliseconds). In this simulation, each sender is in either the “on” (sending) state or the “off” (idle) state; the durations of each state are drawn from a Pareto distribution (which has a “heavy tail”).

3. it can provide variable data rates to different communications essentially on an “as needed” basis.

At the same time, notice that because there is no reservation of resources, packets could arrive faster than can be sent over a link, and the switch must be able to handle such situations. Switches deal with transient bursts of traffic that arrive faster than a link’s bit rate using *queues*. We will spend some time understanding what a queue does and how it absorbs bursts, but for now, let’s assume that a switch has large queues and understand why packet switching actually works.

Packet switching supports end points sending data at variable rates. If a large number of end points conspired to send data in a synchronized way to exercise a link at the same time, then one would end up having to provision a link to handle the peak synchronized rate for packet switching to provide reasonable service to all the concurrent communications.

Fortunately, at least in a network with benign, or even greedy individual communicating pairs, it is highly unlikely that all the senders will be perfectly synchronized. Even when senders send long bursts of traffic, as long as they alternate between “on” and “off” states and move between these states at random (the probability distributions for these could be complicated and involve “heavy tails” and high variances), the aggregate traffic




Figure 15-7: Network traffic variability.

of multiple senders tends to smooth out a bit.⁷

An example is shown in Figure 15-6. The x-axis is time in milliseconds and the y-axis shows the bit rate of the set of senders. Each sender has a link with a fixed bit rate connecting it to the switch. The picture shows how the aggregate bit rate over this short time-scale (4 seconds), though variable, becomes smoother as more senders share the link. This kind of multiplexing relies on the randomness inherent in the concurrent communications, and is called *statistical multiplexing*.

Real-world traffic has bigger bursts than shown in this picture and the data rate usually varies by a large amount depending on time of day. Figure 15-7 shows the bit rates observed at an MIT lab for different network applications. Each point on the y-axis is a 5-minute average, so it doesn't show the variations over smaller time-scales as in the previous figure. However, it shows how much variation there is with time-of-day.

So far, we have discussed how the aggregation of multiple sources sending data tends to smooth out traffic a bit, enabling the network designer to avoid provisioning a link for the sum of the peak offered loads of the sources. In addition, for the packet switching idea to really work, one needs to appreciate the time-scales over which bursts of traffic occur in real life.

What better example to use than traffic generated over the duration of a 6.02 lecture on the 802.11 wireless LAN in 34-101 to illustrate the point?! We captured all the traffic that traversed this shared wireless network on a few days during lecture. On a typical day, about a Gigabyte of traffic goes through it over a 50-minute lecture, with numerous applications in the mix. As one might expect, most of the traffic is due to Bittorrent, Web browsing, email, with the occasional IM sessions thrown in the mix. Domain name system (DNS) lookups, which are used by most Internet applications, also generate a sizable number of packets (but not bytes).

⁷It's worth noting that many large-scale *distributed denial-of-service attacks* try to take out web sites by saturating its link with a huge number of synchronized requests or garbage packets, each of which individually takes up only a tiny fraction of the link.




Figure 15-8: Traffic bursts at different time-scales, showing some smoothing. Bursts still persist, though.

Figure 15-8 shows the aggregate amount of data, in bytes, as a function of time, over different time durations. The top picture shows the data over 10 millisecond windows— here, each y-axis point is the total number of bytes observed over the wireless network corresponding to a non-overlapping 10-millisecond time window. We show the data here for a randomly chosen time period that lasts 17 seconds. The most noteworthy aspect of this picture is the bursts that are evident: the maximum (not shown is as high as 50 Kbytes over this duration, but also note how successive time windows could change between close to 20 Kbytes and nearly 0. From time to time, larger bursts occur where the network is essentially continuously in use (for example, starting at 14:12:38.55).

The middle picture shows what happens when we look at 100 millisecond windows. Clearly, bursts persist, but the variance has reduced. When we move to 1 second windows, we see the same effect persisting, though again it's worth noting that the bursts don't actually disappear.

These data sets exemplify the traffic dynamics that a network designer has to plan for while designing a network. One could pick a data rate that is higher than the peak expected over a short time-scale, but that would be several times larger than picking a smaller value and using a queue to absorb the bursts and send out packets over a link of a smaller rate. In practice, this problem is complicated because network sources are not




Figure 15-9: Packet switching uses queues to buffer bursts of packets that have arrived at a rate faster than the bit rate of the link.

“open loop”, but actually react to how the network responds to previously sent traffic. Understanding how this feedback system works is beyond the scope of 6.02; here, we will look at how queues work.

15.3.2 Absorbing bursts with queues

Queues are a crucial component in any packet switch. They absorb bursts of data (see Figure 15-9, so a key question is how big they should be. The reason this question is important is that if you make them too big, all that happens is that delays grow (for queues by themselves don’t make data transfers go any faster), but if you make them too small, the switch may drop too many packets. We won’t be able to completely address this problem in this course, but it is a question that we will discuss further in a later lecture.

Queues add delay and these delays are variable. In fact, queues are only one of four different sources of delays in networks. The next section describes these sources and talks in more detail about queuing delays and how to analyze them. The key result that we will look at here is called Little’s law.

15.4 Understanding Network Delays

An important measure of performance of any communication network is the network delay, also called the network latency. There are four sources of delay:

1. **Propagation delay.** This source of delay is due to the fundamental limit on the time it takes to send any signal over the medium. For a wire, it’s the speed of light over that material (for typical fiber links, it’s about two-thirds the speed of light in vacuum). For radio communication, it’s the speed of light in vacuum (air), about 3×10^8 meters/second.

The best way to think about the propagation delay for a link is that it is equal to the time for the first bit of any transmission to reach the intended destination. For a path comprising multiple links, just add up the individual propagation delays.

2. **Processing delay.** Whenever a packet (or data frame) enters a switch, it needs to be processed before it is sent over the outgoing link. In a packet-switched network, this processing involves, at the least, looking up the header of the packet in a table to determine the outgoing link. It may also involve operations like computing a packet checksum, modifying the packet's header, etc. The total time taken for all such operations is called the processing delay of the switch.
3. **Transmission delay.** The transmission delay of a link is the time it takes for a packet of size S bits to traverse the link. If the bit rate of the link is R bits/second, then the transmission delay is S/R seconds.
4. **Queueing delay.** Queues are a fundamental data structure used in packet-switched networks to absorb bursts of data arriving for an outgoing link at speeds that are (transiently) faster than the link's bit rate. The time spent by a packet *waiting* in the queue is its queueing delay.

Unlike the other components mentioned above, the queueing delay is usually variable. In many networks, it might also be the dominant source of delay, accounting for about 50% (or more) of the delay experienced by packets when the network is congested. In some networks, such as those with satellite links, the propagation delay could be the dominant source of delay.

15.4.1 Little's Law

A common method used by engineers to analyze network performance, particularly delay and throughput (the rate at which packets are delivered), is *queueing theory*. In this course, we will use an important, widely applicable result from queueing theory, called *Little's law* (or Little's theorem).⁸ It's used widely in the performance evaluation of systems ranging from communication networks to factory floors to manufacturing systems.

For any stable (i.e., where the queues aren't growing without bound) queueing system, Little's law relates the average arrival rate of items (e.g., packets), λ , the average delay experienced by an item in the queue, D , and the average number of items in the queue, N . The formula is simple and intuitive:

$$N = \lambda \times D \quad (15.1)$$

Example. Suppose packets arrive at an average rate of 1000 packets per second into a switch, and the rate of the outgoing link is larger than this number. (If the outgoing rate is smaller, then the queue will grow unbounded.) It doesn't matter how inter-packet arrivals are distributed; packets could arrive in weird bursts according to complicated distributions. Now, suppose there are 50 packets in the queue on average. That is, if we sample the queue size at random points in time and take the average, the number is 50 packets.

Then, from Little's law, we can conclude that **the average queueing delay experienced by a packet is 50/1000 seconds = 50 milliseconds.**

⁸This "queueing formula" was first proved in a general setting by John D.C. Little, who is now an Institute Professor at MIT (he also received his PhD from MIT in 1955). In addition to the result that bears his name, he is a pioneer in marketing science.




Figure 15-10: Packet arrivals into a queue, illustrating Little's law.

Little's law is quite remarkable because it is independent of how items (packets) arrive or are serviced by the queue. Packets could arrive according to any distribution. They can be serviced in any order, not just FIFO. They can be of any size. In fact, about the only practical requirement is that the queueing system be stable. It's a useful result that can be used profitably in back-of-the-envelope calculations to assess the performance of real systems.

Why does this result hold? Proving the result in its full generality is beyond the scope of this course, but we can show it quite easily with a few simplifying assumptions using an essentially pictorial argument. The argument is instructive and sheds some light into the dynamics of packets in a queue.

Figure 15-10 shows $n(t)$, the number of packets in a queue, as a function of time t . Each time a packet enters the queue, $n(t)$ increases by 1. Each time the packet leaves, $n(t)$ decreases by 1. The result is the step-wise curve like the one shown in the picture.

For simplicity, we will assume that the queue size is 0 at time 0 and that there is some time $T \gg 0$ at which the queue empties to 0. We will also assume that the queue services jobs in FIFO order (note that the formula holds whether these assumptions are true or not). Let P be the total number of packets forwarded by the switch in time T (obviously, in our special case when the queue fully empties, this number is the same as the number that

entered the system).

Now, we need to define N , λ , and D . One can think of N as the *time average* of the number of packets in the queue; i.e.,

$$N = \frac{1}{T} \int_{t=0}^T n(t) dt.$$

The rate λ is simply equal to P/T , for the system processed P packets in time T .

D , the average delay, can be calculated with a little trick. Imagine taking the total area under the $n(t)$ curve and assigning it to packets as shown in Figure 15-10. That is, packets A, B, C, ... each are assigned the different rectangles shown. The height of each rectangle is 1 (i.e., one packet) and the length is the time until some packet leaves the system. Each packet's rectangle(s) last until the packet itself leaves the system.

Now, it should be clear that the time spent by any given packet is just the sum of the areas of the rectangles labeled by that packet.

Therefore, the average delay experienced by a packet, D , is simply the area under the $n(t)$ curve divided by the number of packets. That's because the total area under the curve, which is $n(t)$ is total delay experienced by all packets.

Hence,

$$D = \frac{\int_{t=0}^T n(t) dt}{P}.$$

From the above expressions, Little's law follows: $N = \lambda \times D$.

Problems and Questions

These questions are to help you improve your understanding of the concepts discussed in this lecture. The ones marked ***PSet*** are in the online problem set. Some of these problems will be discussed in recitation sections. If you need help with any of these questions, please ask anyone on the staff.

1. Under what conditions would circuit switching be a better network design than packet switching?
2. Which of these statements are correct?
 - (a) Switches in a circuit-switched network process connection establishment and tear-down messages, whereas switches in a packet-switched network do not.
 - (b) Under some circumstances, a circuit-switched network may prevent some senders from starting new conversations.
 - (c) Once a connection is correctly established, a switch in a circuit-switched network can forward data correctly without requiring data frames to include a destination address.
 - (d) Unlike in packet switching, switches in circuit-switched networks *do not* need any information about the network topology to function correctly.
3. Consider a switch that uses time division multiplexing (rather than statistical multiplexing) to share a link between four concurrent connections (A, B, C, and D) whose packets arrive in bursts. The link's data rate is 1 packet per time slot. Assume that the switch runs for a very long time.
 - (a) The average packet arrival rates of the four connections (A through D), in packets per time slot, are 0.2, 0.2, 0.1, and 0.1 respectively. The average delays observed at the switch (in time slots) are 10, 10, 5, and 5. What are the average queue lengths of the four queues (A through D) at the switch?
 - (b) Connection A's packet arrival rate now changes to 0.4 packets per time slot. All the other connections have the same arrival rates and the switch runs unchanged. What are the average queue lengths of the four queues (A through D) now?

4. ***PSet*** Over many months, you and your friends have painstakingly collected a 1,000 Gigabytes (aka 1 Terabyte) worth of movies on computers in your dorm (we won't ask where the movies came from). To avoid losing it, you'd like to back the data up on to a computer belonging to one of your friends in New York.

You have two options:

- A. Send the data over the Internet to the computer in New York. The data rate for transmitting information across the Internet from your dorm to New York is 1 Megabyte per second.
- B. Copy the data over to a set of disks, which you can do at 100 Megabytes per second (thank you, firewire!). Then rely on the US Postal Service to send the disks by mail, which takes 7 days.

Which of these two options (A or B) is faster? And by how much? Note on units:

1 kilobyte = 10^3 bytes

1 megabyte = 1000 kilobytes = 10^6 bytes

1 gigabyte = 1000 megabytes = 10^9 bytes


1 terabyte = 1000 gigabytes = 10^{12} bytes

5. Little's law can be applied to a variety of problems in other fields. Here are some simple examples for you to work out.
- (a) F freshmen enter MIT every year on average. Some leave after their SB degrees (four years), the rest leave after their MEng (five years). No one drops out (yes, really). The total number of SB and MEng students at MIT is N .
What fraction of students do an MEng?
 - (b) A hardware vendor manufactures \$300 million worth of equipment per year. On average, the company has \$45 million in accounts receivable. How much time elapses between invoicing and payment?
 - (c) While reading a newspaper, you come across a sentence claiming that "*less than 1% of the people in the world die every year*". Using Little's law (and some common sense!), explain whether you would agree or disagree with this claim. Assume that the number of people in the world does not decrease during the year (this assumption holds).
6. You send a stream of packets of size 1000 bytes each across a network path from Cambridge to Berkeley. You find that the one-way delay varies between 50 ms (in the absence of any queueing) and 125 ms (full queue), with an average of 75 ms. The transmission rate at the sender is 1 Mbit/s; the receiver gets packets at the same rate without any packet loss.
- A. What is the mean number of packets in the queue at the bottleneck link along the path (assume that any queueing happens at just one switch).

You now increase the transmission rate to 2 Mbits/s. You find that the receiver gets packets at a rate of 1.6 Mbits/s. The average queue length does not change appreciably from before.

- B. What is the packet loss rate at the switch?
- C. What is the average one-way delay now?

7. Consider the network topology shown below.



- (a) The sender sends packets back-to-back with a negligible inter-packet delay. The queue has no other packets. What is the time delay between the arrival of the first bit of each packet at the receiver?
 - (b) The receiver acknowledges each packet to the sender, and each acknowledgment has a size $A = 100$ bytes. What is the minimum possible round trip time between the sender and receiver? The round trip time is defined as the duration between the transmission of a packet and the receipt of an acknowledgment for it.
8. ***PSet*** The wireless network provider at a hotel wants to make sure that anyone trying to access the network is properly authorized and their credit card charged before being allowed. This billing system has the following property: if the average number of requests currently being processed is N , then the average delay for the request is $a + bN^2$, where a and b are constants. What is the maximum rate (in requests per second) at which the billing server can serve requests?

Acknowledgments

Many thanks to Sari Canelake and Lavanya Sharan for their careful reading and helpful comments.