

CMPT 165: Review

Tamara Smyth, tamaras@cs.sfu.ca
School of Computing Science,
Simon Fraser University

December 5, 2011

Access to the Web

- Access to the Web requires:
 - a computer (of some kind)
 - a connection to the Internet
 - special software on the computer
 - * Internet communication software (invisible to the user)
 - * application software (visible to the user), called a *browser*.

Coverage

- The usefulness of a technology is determined, in part, by its **coverage**:

If no one else has a communication service, it's useless; if everyone else has the service, it becomes a necessity.

- Coverage ranges between
 - local service
 - * not available to everyone
 - * limited usefulness
 - universal service
 - * available to everyone
 - * limited usefulness
- Coverage is affected by
 - availability: to whom is the technology available?
 - affordability: what is the initial and subsequent financial risk?

Basic Communication Technology

- Analog to Digital conversion
- Modulation/Demodulation
- Sharing
- Packet Switching (sharing a transmission path).
- Routers
- Connecting to the Internet

Analog Signals

- The term “analog” refers to the fact that a signal is “analogous” of the signal it represents.




Figure 1: An *analog* electrical signal represents the pressure variations of a sound wave.

- Microphone and speaker's are analog devices because they record and play an exact analog of sound:
- Analog signals are *continuous* in time.
- A continuous signal cannot be stored, or processed, in a computer since it would require an infinite amount of data.
- Analog signals must therefore be *discretized*, or *digitized*, to produce a finite set of numbers, for computer use.

Digital Signals

- The process of taking individual values on a continuous signal at regular time intervals is called *sampling*.
- Music on compact discs is stored digitally. The sampling rate is typically 44,100 samples per second (or 44.1 kHz).




Figure 2: A digital, or sampled, signal.

Representing a Digital Signal

- Computers represent data (sound, music, text, numbers, video, image) or operating instructions in *binary form*.
- The binary system is based on the number 2; each bit position in the group is worth 2 times as much as the position to the right.

- The binary number 1001 is evaluated as

$$1 \times 2^3 + 0 \times 2^2 + 0 \times 2^1 + 1 \times 2^0 = 9$$

- Other systems include octal and hexadecimal.

Bits

- The smallest unit of memory is the *bit*, representing the position, either “on” or “off”, of an electronic switch inside the computer hardware.
- A bit can represent either “0” or “1” depending on the state of the switch.
- The term “bit” is a contraction of the term “binary digit”.
- Two bits can assume 4 unique patterns/values:

Binary form	Decimal form
00	0
01	1
10	2
11	3

- In general, n bits can assume 2^n unique values.
- The *byte* is nearly always defined as 8 bits and can therefore assume $2^8 = 256$ unique values.

Error Checking

- Interference from natural phenomena, distortion from electrical or magnetic interference can cause damage or loss of data.
- A **parity bit** is an extra bit added to a code before transmission to help the hardware detect errors. It is set to
 - 1: if a sequence has an *odd* number of 1 bits
 - 0: if a sequence has an *even* number of 1 bits.
- **Checksum** is a technique for detecting damaged bits. The software sends the sum of the numbers in the message. If a message consists of the numbers

1, 3, 5

the software sends

1, 3, 5, 9

- When the messages arrives, the receiving software sums all the values except the last, and then checks to ensure the sum equals the last value.

Modulation

- Modulation was discovered as a means of transmitting voice (and other signals).
- To transmit a signal, the sender must have a device called a *modulator*, which uses 2 signals:
 - **carrier**: a basic electrical signal that oscillates back and forth; a high-frequency periodic waveform
 - **modulator**: a second signal containing the information being transmitted, which “modulates” or changes the carrier slightly.




Figure 3: An example of amplitude modulation.

Modem

- The receiver must have a device called a demodulator, which performs the reverse function known as “demodulation”.
- The demodulator is tuned to the expected frequency and recovers the modulator, by measuring how much the incoming signal deviates from the expected carrier frequency.
- A *modem* is a modulator and a demodulator.
- Before two computers in the Internet can communicate, a modem must be installed at each end of the connection.

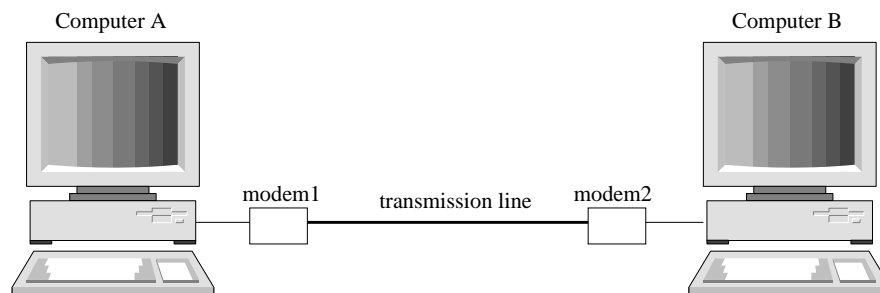


Figure 4: Two computers/ modems connected by a transmission line.

- Modems allow for data to be sent between them in both directions.

Connecting Computers

- How/why do we connect computers?
- LAN technologies were motivated by 2 trends:
 1. computers had grown smaller and less expensive (goodbye expensive mainframes!)
 2. computers helping with ordinary office tasks
- Low-cost computers meant most organizations could now have several computers.
- Data transfer would usually involve removable media storage devices (magnetic tapes or disks).

Sharing

- In a network, there isn't a single wire connecting each communicating computer. Transmission paths must be shared.
- Only one data transfer can occur on a given wire at a time.
- What happens when one customer's order is very large?

Granting one party exclusive access of a shared transport path can delay all other parties.

Preventing long delays

- How do we prevent delays and long wait times?
 - **Channels:** use signal modulation on multiple channels, mix them, and transmit across a cable (e.g. cable company).
 - **Taking Turns:** allow one computer access to the network at a time and limit the amount of data that a computer can transfer on each turn.
- Though it is possible to use both, most network technologies use the second idea, called *packet switching*.

Packets

- The unit of data that can be transferred at one time over a network is called a *packet*.
- The packet *header* includes:
 1. the origin address
 2. the destination address
 3. the position of the packet in the message (so it can be reassembled at the destination in the correct order).
- Though the packet *size* is limited, it may contain any amount of data up to the limit.

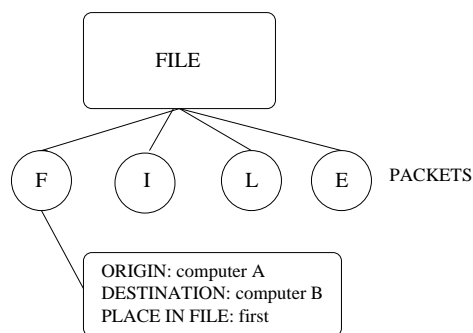


Figure 5: All data is transfer across the Internet in packets. A sender divides a message or document into packets.

Packet Switching

- A message is divided into as many packets as are needed to accommodate the message size.

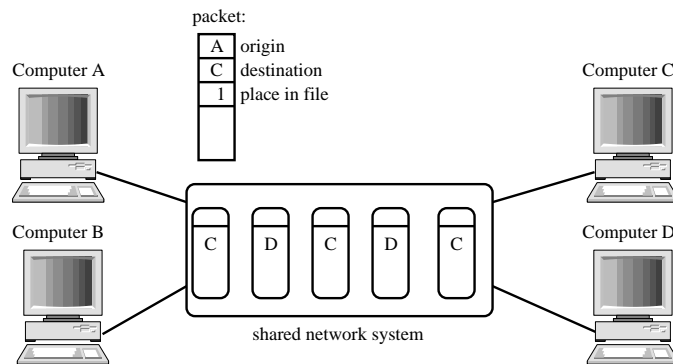


Figure 6: Computer A is sharing a transmission path with Computer B by taking turns sending packets.

- One computer will not have to wait for another computer's entire message to be transmitted before it has the opportunity to send packets.

Packet switching allows many communications to proceed simultaneously without requiring an application to wait for all other communication to complete.

Who uses packet switching?

- Packet switching can be used by many devices: printers, scanners, video cameras, etc.
- Both LANs and WANs use packet switching.
- The Internet is a packet switching system—packets from many machines traverse the Internet at the same time.
- A receiver reassembles the original message from the packets that arrive.

Extending a connection to a network

- A computer's (usually short) connection to a LAN may be extended using:
 1. two modems with wires connecting them

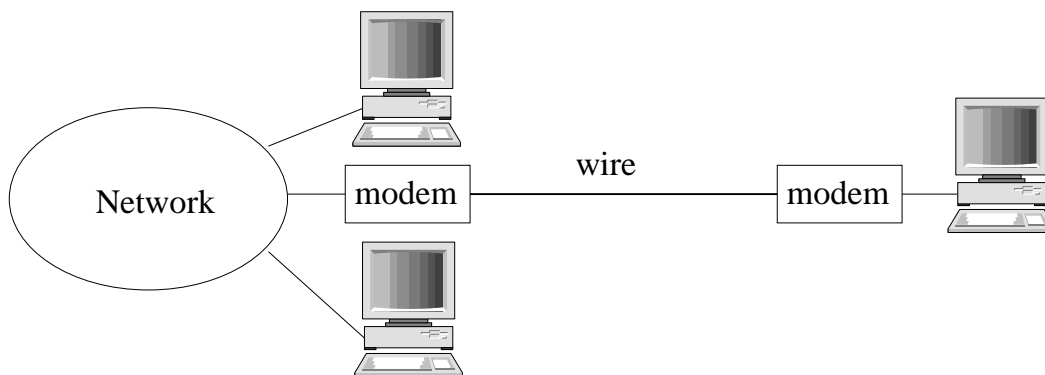


Figure 7: Two modems connected by a wire

2. two modems connected via the telephone system

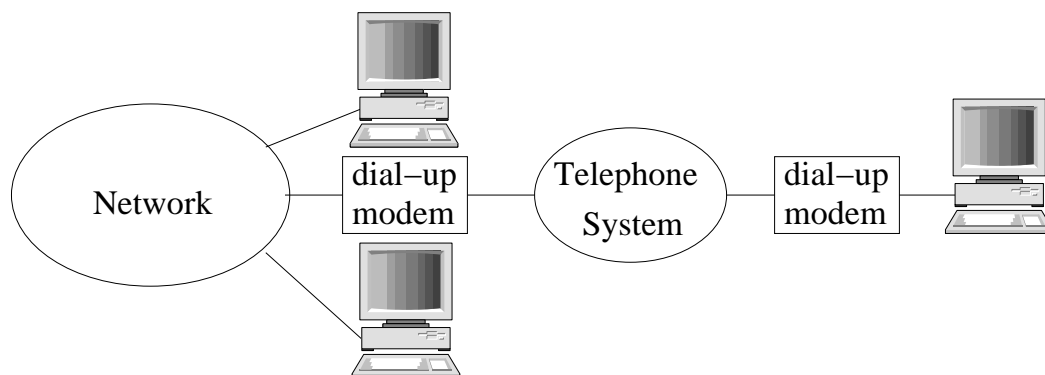


Figure 8: Two modems connected by the telephone system

3. two optical modems connected by glass fiber:




Figure 9: Two modems connected by a fiber optic cable

Interconnecting networks—Routers

- Expanding networks it not necessarily straight forward. Several packet switching technologies exist and they all are not compatible.
- A computer can have multiple circuit boards connecting it to several networks (and network technologies) simultaneously.
- The Internet uses such a special-purpose computers to interconnect networks called *routers*.

Routing and Packet Forwarding

- The router has one major task: forward packets from one network to the other.
- The router receives a packet sent to it across one network, and sends the packet on to its destination across another network, a process called *packet forwarding*




Figure 10: A computer on Network 1 can send packets to a computer on a Network 2 via a router. The networks need *not* be the same type.

The Internet: a network of networks

- The Internet is not a conventional network—it consists of thousands of computer networks, interconnected by routers.
- A router is a fundamental Internet technology; because of it, the Internet can contain multiple types of networks.
- *Backbone* network is a term used to describe a major WAN to which other networks attach.

How do individuals connect to the Internet?

- Routers allow multiple networks to connect to the Internet, but how do individual computers connect?
- A company that provides Internet access is known as an Internet Service Provider (ISP). They will charge for
 1. using the internet
 2. a physical connection to the Internet (a separate dedicated connection between a site and an ISP).
- ISPs face conflicting goals: trying to create a technology that performs well (so customers don't have to wait) but at low cost.
- **Last Mile Problem:** the final leg of delivering connectivity from an ISP to a customer. It is typically seen as expensive because “fanning out” wires and cable is a considerable physical undertaking.

Early Internet Connections

- Until mid-1990s...
- **leased digital circuit:** only available technology offering a *dedicated connection* to the Internet
 - Pros: availability and control (the only traffic is that of the customer)
 - Cons: *very expensive*, unavailable to the individual
- **dial-up:** access through the telephone system: a computer uses a modem and software to connect; software instructs the modem to place a telephone call to the number provided by the ISP.
 - Pros: affordable and available to the individual
 - Cons: slower; overhead in making a connection

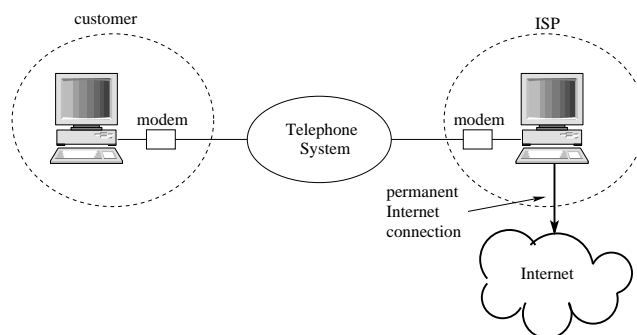


Figure 11: A computer connected to the Internet with a dial-up connection.

Continuous connectivity

- A dedicated (leased) circuit offered huge advantages:
 1. higher speed data delivery
 2. instantaneous access and continuous availability (no waiting for the telephone to be free before using the Internet).

Continuous connectivity has effected the way people use the Internet. Access is no longer a time-consuming activity.

Continuous Connectivity to Individuals

- During 1990s, ISPs began offering:
 1. **Cable modem:**
 - uses the same wiring as cable television
 - cable modified for digital bi-directional signals
 - shared wiring could reduce user's bandwidth
 2. **DSL: Asymmetric Digital Subscriber Line (ADSL)**
 - uses the same wiring as a conventional telephone
 - bandwidth not shared as with cable modem
 - distance limitations: difficult to reach rural areas
 3. **Wireless access:**
 - use radio transmission (like cell phone systems)
 - helps solve the last mile problem: a transmitter is deployed in a geographic region (often on a utility pole) and in the customer's home
- All designed so that data transmission does not interfere with simultaneous television/cable signals.
- All provide speed, continuous and instantaneous connectivity, affordability, and use of existing wiring.

Communication Protocol

- As with humans, for two computers to communicate they must share a common language.

A communication protocol is an agreement that specifies a common language two computers use to exchange messages.

- Communication is defined by specifying:
 - the format and meaning of a computer's message
 - conditions under which a computer should send a message
 - how a computer should respond when a message arrives.

Internet Protocol (IP)

- Internet Protocol (IP) is the key communication protocol used in the Internet specifying:
 - the rules that define how computers communicate/connect.
 - how a packet must be formatted
 - how a router must forward each packet toward its destination

IP Software

- Computer hardware does not understand IP—it requires IP software.
- Because IP is fundamental, each computer usually has a single copy of IP software that all applications share.
- **IP Datagrams:**
 - Each packet sent across the Internet must follow the format specified by the IP and are called IP datagrams (to distinguish from packets for other networks).
 - IP Datagrams are not limited to the packed format used by the underlying (nor any particular) network technology

Any computer connected to the Internet with IP software can create an IP datagram and send it to any other computer on the Internet.

The Internet: A network of networks

- IP transforms a collection of networks and routers into a seamless communication system by making the Internet function like a single, large network.




Figure 12: Computers connect to a *seemingly* single large network.

- Users remain unaware of the Internet's networks and routers, just as telephone subscribers remain oblivious to the wires and switches that comprise the telephone system.

The Internet Reality

- The Internet actually contains a complex internal physical structure that users never see.




Figure 13: Datagrams travelling between computers uses a physical path: either a network, or through a router to another network.

- A datagram travelling across the Internet must follow a physical path from one computer to another
- Each network technology defines its own packet format and does not understand IP datagrams.

Handling IP Datagrams

- The internet sends an IP datagram across a single network by placing it inside a network packet, making it appear as simply data.
- When the network packet arrives at the next computer, the computer “opens” the packet and extracts the datagram to determine how to process it.
- When a router determines that the datagram must be sent across another network, the router creates a new network packet.

IP Addresses (Syntax)

- Recall that computers (including routers) connected to the Internet are assigned a unique *IP address*.
- One computer must know the address of another computer before it can communicate with it.
- An IP address is stored in four bytes (or 32 bits) which means there are only 2^{32} possible addresses
- When displayed by an application it is written as 4 decimal numbers, such as: 128.10.2.1
- Mostly however, humans use names to specify a computer rather than a numeric address.

Communicating on different networks

- Example: C1 communicates with C3 on a different network:
 - IP software on C1 must create an IP datagram: header specifies the sender (C1) and receiver's (C3) IP address
 - Since C1 and C3 are on different networks, the IP datagram must pass through a router
 - The router examines the destination address to determine where to send it
 - The router sends it either across a network to another router (which repeats the above step) or if it can, across the network to C3.

Network Congestion and Delivery Error

- If two networks with the same capacity, A and B respectively, are connected to a third network D with the same capacity, packets from A and B cannot flow to network D at full speed.
- When cars from two lanes merge into a single lane, the traffic slows down or stops completely (traffic jam!).
- On the Internet, datagrams are discarded!
- Though routers have memory and can store data during temporary congestion, if datagrams continue to arrive faster than they can leave, the router discards them until congestion clears.
- Packet switching systems, like those used in the Internet, need additional communication software to ensure that data is delivered
- Transmission Control Protocol (TCP) is the second major communication protocol (after IP)

TCP helps IP guarantee delivery

- Because packet switching hardware can become overrun with datagrams, TCP is needed to improve the Internet's *reliability*.
- As most computers attached to the Internet run IP software, most also run TCP software.
- Both work so well together, and are so important, the entire set of communication protocols the Internet uses is known as TCP/IP protocol suite.
- TCP checks for lost datagrams.
- Since there are usually multiple paths datagrams can take (“detouring around failed routers or networks that are down), they don't necessarily arrive at their destination in the right order.
- TCP automatically checks incoming datagrams and puts the data back in order.
- TCP also checks for duplicate datagrams and only accepts the first copy to arrive.

Handling lost datagrams

- TCP handles lost datagram by using *timers* and *acknowledgements*.
- When data arrives at its final destination, the receiver's TCP software sends an acknowledgment back to the source to confirm delivery.
- Whenever TCP software sends data, it starts a timer using the computer's internal clock:
 - if an acknowledgement arrives before the timer expires, TCP cancels the timer
 - if the timer expires before an acknowledgment arrives, TCP assumes the datagram was lost and resends.
- TCP's strategy is adaptable:
 - if the destination is close to the source, TCP only waits a short time before resending, if it is far, TCP waits longer.
 - TCP measures current delays on the Internet and adjusts the timeout mechanism.

TCP/IP

- Although TCP and IP can be used separately, they were designed at the same time to work as part of a unified system, without duplicating the work done by each.
- A computer connected to the Internet needs BOTH TCP and IP software. IP provides a way to transfer a packet from its source to its destination.
- TCP handles problems like datagram loss or delivery out of order. Together they provide a reliable way to send data across the Internet.”
- Vendors often sell a single package that includes software for TCP, IP and a few related communication protocols. Collectively, the set is known as TCP/IP software.