Introduction and History of Modems

Introduction

The modem is a device used in affecting communications between computers across telephone lines. The word modem is derived from the combination of the words modulation and demodulation. For it is these two functions that the modem performs. It takes digital data from a computer and modulates it into an analog signal to be passed along the telephone line. At the other end, another modem demodulates the signal back to the original digital form so the receiving computer can interpret it.

History

It was in the 1950s that the first modems were being developed. There was a need to transmit data for North American air defense, so efforts were made to accomplish the goal of data transfer across the existing telephone wires. The air defense was using modems by the end of the 1950s, but the first commercial device was not available until 1962. It was called the Bell 103, by AT&T. This first modem allowed full-duplex transmission, and boasted data rates up to 300 bits per second. Shortly after the Bell 103, there came the Bell 212, which reached speeds of 1200 bits per second. It also employed a method of modulation called phase-shift keying (PSK). This was a step up from the frequency-shift keying (FSK) method that the Bell 103 employed.

Over the next fifteen years, the efforts were to make the modems transmit data at a higher rate. In order to accomplish this, the telephone system required some improvement. As it was, due to mutual interference of signals being attenuated at various rates through the system, there was smearing of data symbols. To compensate for this, equalizers needed to be applied to the telephone lines. The automatic adaptive equalizer was invented in 1965 at Bell Laboratories by Robert Lucky. While equalizers had been used for some time, they required human intervention to be adjusted appropriately. With the advent of the automatic adaptive equalizer, data could be transmitted at high rates, as was desired. Modem technology also improved in this time, and by 1980, there existed modems that could transmit up to 14.4 kilobits per second over four-wire leased lines.

By 1984, modems were to the point of transmitting 9.6 kilobits per second over a single-pair circuit on the telephone system. To make this a reality, advances were made in echo cancellation, which keeps the sending modem from picking up its transmitted signal on its own receiver. This problem, of course, only presented itself when trying to send high speed data over a single circuit. Additionally, a new coded modulation with error correcting codes was developed. This integral error correction made the signal less susceptible to noise.

Using the same sort of technology, modem speeds were increased to 14.4 kilobits per second by 1991. Then, in 1994, it doubled to 28.8 kilobits per second. Soon after, there came 33.6, which was thought to be an upper limit for phone line transmissions. But along came the 56k modem, and a new set of standards, so the speeds continue to push the envelope of the capacity of the telephone system.

Physical Layer of Modems

Basically, a modem consists of a power supply, transmitter, and receiver. The power supply provides the voltage necessary to run the modem's circuitry. Depending on the type, the power may be supplied from an external source, or through a connection with the computer's power supply. The transmitter portion contains a modulator, which converts the digital data into an analog signal. Also in the transmitter is circuitry that filters, wave shapes, and signal controls while the signal is being modulated. The receiver had a demodulator and associated circuitry that reverses the modulation process and returns the original digital signal.

The modem connects to the Public Switched Telephone Network (PSTN) through a standard twisted pair phone cable connected to a phone jack. Modems can be either internal or external to the CPU of a computer. Sometimes the modem is attached directly to the logic board. Other times, the connection is made externally through a port. The port configuration is based on the RS-232-C standard, where the connector consists of 25 pins. Frequently, only a subset of the standard pins is actually built into the connectors. An 8- or 9-pin connector will work just fine, simply using the pins for the more frequent commands. As long as there is a wire each for sending, receiving, and ground, the connector will function.

Multiplexing

Multiplexing the method of sharing a channel among users. The phone system, and thus the modem system, uses Time Division Multiplexing. This method of multiplexing give specific time slots to each user, and that user is only allowed to transmit during that time.

Modulation

In order to convert a signal from digital to analog, then back to digital (which is what modems do), a method of modulation must be used. To convert from digital to analog, the most efficient form of multiplexing available is Quadature Amplitude Modulation, which is a combination of frequencies, amplitudes, and phase shifts. This means that many different signals can be represented with a small set of frequencies, amplitudes, and phase shifts. The frequencies and other shifts need to be widely dispersed, so that it is less likely for two signals to be confused.

Most modern modems use this form of modulation. Older ones use Frequency modulation, which of course, only changes the frequency.

Network Topology

The topology of the telephone system is switched, or star. This topology has a central computer that all other devices are connected to. It is through this central computer that all communication occurs. The phone system is actually many separate stars that are connected together into a larger network, but the concept is the same.

Data Link Layer of Modems

Data Transparency

A modem uses the bit-oriented protocol for its frames. So bit stuffing is required to differentiate the data being transmitted from the flag bit sequence. The bit stuffing imposes an overhead of up to 20%. This is much better than byte stuffing, which could double the number of bytes to be transmitted.

Data Transmission

All modern modems are full-duplex (since the V.21 standard).

Error Detection/Correction

Today's modems use a Cyclical Redundancy Check (CRC) for error detection. In this check, an N must be chosen, and that is used for the polynomial checking sequence. The CRC will detect all single-bit errors, most double-bit errors, and most errors where the number of bits in error is odd. CRC will also detect all burst errors where the burst length is less than N. It will detect most where the burst length is greater than or equal to N. Thus, in order to detect as many errors as possible, the N should be chosen to be larger than the expected burst length.

Once errors are detected, they need to be corrected. The phone system uses a Selective Repeat Request. In this case, the information frames are sent to the receiver with a sequence number attached. This lets the receiver know what frames it should have received so far. If the receiver gets a frame out of order, it will send notification to the sender about that, and the sender will retransmit that frame only. The receiver also sends acknowledgements for frames receives, so in the case where the sender transmits a frame and receives nothing back, it will automatically retransmit after a certain time has elapsed. This assures that all the data that is meant to be transmitted actually arrives successfully.

Flow Control

Some modems use the X-ON/X-OFF flow control protocol. In this protocol, an X-ON control character (DC1 -- ASCII code 17) is transmitted. When received, the recipient knows that it can send data. Data continues to be transmitted until an X-OFF (DC3 -- ASCII code 19) is received. That tells the station sending data that the one receiving data is filling up, and can't take any more data. So that transmission stops until another X-ON is received. Of course, this flow control occurs the same in both directions, since it is a full-duplex system.

However, putting control characters in the data stream can cause problems. The control characters may appear in binary data, which will stop the data transfer unexpectedly. To avoid this, a flow control called RTS/CTS (Request to Send/Clear to Send). These are pins that are on the connector and signals are sent through these when the transmission should be stopped or started.

Data Compression/Coding

The only type of data compression that the modem itself can do is lossless compression. This type of compression allows the original data to be reconstructed perfectly once it is received. It is most successful for text-based information, since most binary information, especially images, sounds, movies, etc. have already had lossy compression applied to them, so further compression for documents of that type is not really possible.

Modems use ASCII character encoding. In order to code the ASCII characters into less than 7 bits, an encoding technique, such as Huffman coding, is used. Huffman coding makes the codes shorter for more common characters (2 or 3 bits long, as opposed to 7 bits).

Modem Standards

Modems are used on many different types of computers, function at different speeds, and travel over different quality of wire. But all the modems also need to be able to talk to each other. That is why standards exist. The ITU-T standards are used throughout the world as the model to follow in modem communications. Below is a list of some of these standards.

Standard	Meaning
V.22	Provides 1200 bits per second at 600 baud (state changes per second)
V.22bis	The first true world standard, it allows 2400 bits per second at 600 baud
V.32	Provides 4800 and 9600 bits per second at 2400 baud
V.32bis	Provides 14,400 bits per second or fallback to 12,000, 9600, 7200, and 4800 bits per second
V.32terbo	Provides 19,200 bits per second or fallback to 12,000, 9600, 7200, and 4800 bits per second; can operate at higher data rates with compression; was not a CCITT/ITU standard
V.34	Provides 28,800 bits per second or fallback to 24,000 and 19,200 bits per second and backwards capability with V.32 and V.32bis

V.34bis	Provides up to 33,600 bits per second or fallback to 31,200 or V.34 transfer rates
V.35	The trunk interface between a network access device and a packet network at data rates greater than 19.2 Kbps. V.35 may use the bandwidths of several telephone circuits as a group. There are V.35 Gender Changers and Adapters.
V.42	Same transfer rate as V.32, V.32bis, and other standards but with better error correction and therefore more reliable
V.90	Provides up to 56,000 bits per second downstream (but in practice somewhat less). Derived from the x2 technology of 3Com (US Robotics) and Rockwell's K56flex technology.

The current standard is V.90, and it functions differently from all of its predecessors. This standard was the first to make bit rates of significantly higher than 33.6 kbps a possibility. The way this happened has to do with the way the data is dealt with in the different standards. Prior to V.90, data was always converted to analog to travel over the wire. This is why bit rates were limited. V.90 makes an assumption that one end of the modem connection has a pure digital connection to the phone network. Since this is true for the service providers, the digital-analog conversion is not necessary when downloading data. Data sent from the modem is still sent in analog form, but they are usually keyboard strokes or mouse movements (for web interaction).

Thanks to the phone network being increasingly digital, the physical limit on data transfer increases by a good bit. However, the Federal Communications Commission (FCC) currently restricts the maximum download speeds to 54kbps. So even as technology increases, there is still an imposed limitation.

Case Study: Utilization on a 56k Modem

On a 56k modem, the theoretical bit rate is 53000 bits per second. It is not 56000 bps due to FCC regulations. However, the typical data transfer rate is somewhat slower. This disparity is due to several factors:

- Bandwidth limitations
- Noise on the line
- Digital to Analog conversion
- Analog to Digital conversion
- Framing overhead

Bandwidth

The bandwidth of telephone lines is 3000 Hz. This value directly affects the capacity of the line. Using the Shannon-Hartley law...

- bit rate = bandwidth*(log2(1 + S/N))
- 53000 bps = (3000 Hz)*(log2(1 + S/N))
- so the S/N must be at least 208545. SNR = 10 log(S/N) = 53.2 dB

This high a signal-to-noise ratio is only available on a digital line. The typical SNR on the telephone system is only 35 dB.

As noted in the previous discussion on the V.90 standard, only downloading can be done in this purely digital state. Uploading still goes through the digital-analog-digital conversion, so the actual maximum bit rate in that direction...

- SNR = 35 dB = 10 log(S/N) S/N = 3162.3
- bit rate = (3000 Hz)*(log2(1 + 3162.3) = 34881 bps

Digital-Analog-Digital Conversion

This conversion pretty much only concerns the upload condition, but it deserves treatment here anyway. Say, for argument's sake that the only noise in the channel is due to this conversion between analog and digital. That value is called Quantization noise.

- quantization noise = (1/2^q)^2
- if signal = 1, then S/N = 2^2q
- 2^2q = 3162.3
- q = 5.81 bits

This means that there are 5.81 -> 6 bits needed for this conversion.

Framing

Point-to-Point Protocol (PPP) uses the HDLC frame format. PPP is the framing format used by modem users who want to make more than just a terminal connection to the internet. The HDLC frame format is below:

	Flag	Address	Control	Data	FCS	Flag
# of bits	8	8 or 16	8 or 16	variable	16 or 32	8

According to this, there are a minimum of 48 bits that are not part of the data. These extra bits lower the effective data rate and utilization of the channel.

The optimum frame size in a selective repeat system with a burst error rate of 10^-5 (which are the stats for the phone system) is 277 bytes. That is 2216 bits (F). With the information that there are 48 bits per frame that are not part of the data, that means there are 2168 bits left for data (N). The transmission rate is assumed to be 53000 bps (R) because of the FCC limitation on download data rates. For argument's sake, it takes 1 microsecond to create the frame (T).

- effective data rate = N/(T + (F/R)) = 2168 bits/(.000001 s + (2216 bits/53000 bps)) = 51850 bps
- so using that as a measure of channel utilization = effective data rate/theoretical data rate = 51850 bps/53000 bps = 97.8 %

These, of course, only apply to download rates. In the upload case...

- effective data rate = 2168 bits/(.000001 s + (2216 bits + 6 bits for DAD conversion/34881 bps)) = 34033 bps
- channel utilization = 34033 bps/34881 bps = 97.6 %
- but if we assume the channel capacity is 53000 bps, then utilization = 34033 bps/53000 bps = 64.2 %

So 53000 bps on a 56k modem is only a theoretical capacity. There are several contributing factors that slow the transmission rates down. For instance, high network load will slow down one's individual transmission rates. Excessive noise on the channel can be a factor as well.

Cost Comparison of Modems to Other Services

What follows is a cost comparison of some popular internet services, and their relative speeds. The question being asked here is: are the faster and more costly services actually worth getting?

Note: All prices are current as of 8/1/00

Service	Provider	Set-up fee	Monthly rate	Downstream Data	Upstream Data	Downloading of 1 GB of data, no compression	Total Cost
Various Modem Pool	U of Pitt	N/A	$27.50*	22000 - 44000 bps	unknown	54.23 - 108.46 hours	$81.73 - $135.96
56k	Verizon	$9.95	$19.95	53000 bps	33600 bps	45.02 hours	$74.92
DSL	Verizon	$99.00 (cost of modem)	$39.95	256000 - 640000 bps	90000 bps	3.73 - 9.32 hours	$142.68 - $148.27
Cable Modem	AT&T@Home	no fee	$39.95	"2 seconds on @Home as compared to 9 minutes over 28.8 modem"	128000 bps	.31 hours	$40.26

*This price is based on the $110 per semester Computer Network/Service fee that is charged to a full time student at Pitt.

The calculations do not take actual data throughput into account, but the differences in the calculated times will occur in the same fashion as they do with the theoretical bit rates.

To illustrate the true cost of each of these services, the following values will be used:

- "worth" of an hour = $1
- use the service for one month
- download 1 GB of data each month
- Total cost = download time * $1 per hour + set-up fee + monthly rate

According to these calculations, DSL is actually the most costly overall. Of course, that includes the $99.00 for the DSL modem, which one would presumably only pay for once. If Cable Modem speeds are true to their claim, it is by far the best service in terms of cost. However, since AT&T does not provide any definite specifications on the speed, it is difficult to make an informed decision.

At this point, sticking with the Pitt modem pools is my best choice, especially since I have to pay the network fee anyway.