ASCII Character Codes

	000	001	010	011	100	101	110	111
0000 0	NUL	DLE	SP	0	@	Р	``	р
0001 1	SOH	DC1	!	1	A	Q	а	q
0010 2	STX	DC2	"	2	В	R	b	r
0011 3	ETX	DC3	#	3	С	S	С	S
0100 4	EOT	DC4	\$	4	D	Т	d	t
0101 5	ENQ	NAK	%	5	E	U	е	u
0110 6	ACK	SYN	&	6	F	V	f	V
0111 7	BEL	ETB	,	7	G	W	g	W
1000 8	BS	CAN	(8	Н	Х	h	х
1001 9	HT	EM)	9	I	Y	i	У
1010 A	LF	SUB	*	:	J	Z	j	Z
1011 B	VT	ESC	+	- ,	K	[k	{
1100 C	FF	FS	,	<	L	١	I	
1101 D	CR	GS	-		М]	m	}
1110 E	SO	RS		>	N	^	n	~
1111 F	SI	US	/	?	0	_	0	DEL

COMPSCI 314 S1 C Data Comms - Basics

ASCII Control Codes

(more important ones in red)

NUI Null SOH Start of Header STX Start of Text ETX End of Text FOT ENQ Enquire ACK Acknowledge BFI Bell BS Backspace HT Horizontal Tab IF. Line Feed Vertical Tab VT FF Form Feed CR Carriage Return SO Shift Out SI Shift In

All-zero "filler" character Indicates start of control information Separates heading and text in a message End of text field in a message (or final end, c.f. ETB) End of Transmission *End of transmission with one or more Texts;* alternative meaning is "reset line" Request for response, recipient should reply Positive acknowledgement of msg ("OK to proceed") Activate a visible or audible alarm Move cursor (or print head) back by one position Move cursor (or print head) right to next tab stop Move cursor (or print head) to same position, next line Feed paper to next predefined line Advance to top of next page Move cursor or print mechanism to start of present line Assign special meanings to all following characters End of "Shift Out": all characters revert to their normal meaning

DLE	Data Link Escape	Assign a special meaning to the following character
DC1	Device Control 1	General device control; often "XON" to enable transmission
DC2	Device Control 2	General device control
DC3	Device Control 3	General device control; often "XOFF" to stop transmission
DC4	Device Control 4	General device control
NAK	Negative Acknowledge	Negative response, e.g. bad parity, or device busy
SYN	Synchronous Idle	Establish character synchronisation; also synch. filler code
ETB	End Transmission Block	End of a transmission (more yet to come, c.f. ETX)
CAN	Cancel	Cancel (i.e. ignore) previous data in this message
EM	End of Medium	End of card deck, paper tape, cassette tape etc.
SUB	Substitute	Used to replace a bad or invalid character in message
ESC	Escape	Start of a special character sequence for device control
FS	File Separator	Separates files or other major data units
GS	Group Separator	Separates groups of records within a File
RS	Record Separator	Separator between records
US	Unit Separator	Separator between items in a record

Hierarchical Communication Models

- Communications systems are best handled with a hierarchical control structure, much like a hierarchy or nest of subroutines.
- Some systems, such as TCP/IP, have a very flexible layering, with layers reaching into adjoining layers or even skipping over an adjoining layer.
- We will start with the OSI model (Open Systems Interconnection) which has a very strict seven layer structure — each layer communicates only with its adjoining layers.
- Most later work will be with TCP/IP.
- A half serious comment is that TCP/IP is the way you *do* communications, but OSI is the way you *teach* it!

The OSI Model



- Levels 5 7 are more the concern of MSIS and the designer of applications that use communications.
- We concentrate on levels 1 to 4, and especially levels 2 to 4.
- The details of level 1 are more the concern of Electrical Engineering.

The TCP/IP Protocol suite (to be discussed later)

Users



TELNET	Terminal emulator	FTP	File Transfer Protocol	
SMTP	Simple Mail Transfer Protocol	CMOT	CMIP over TCP	
DNS	Domain Name Server	SNMP	Simple Network Maintenanc	е
			Protocol	
TFTP	Trivial File Transfer Protocol	BOOTP	Boot Protocol	
RPC	Remote Procedure Call	NFS	Network File System	
COMPSCI 314 S	S1 C Data Comms - Basics	21-Mar-05		6

• A "connection" (wire, optical cable, radio, etc.) between two "end-points"

Problem

• Convey a stream of bits or bytes from one end to the other.



• Problems: data coding, transmission, clock recovery, ...

• A Physical Layer able to transmit bits or bytes end to end

Problem

• Convey recognised sequences (messages or frames) of bytes



• Problems: assembling bits into bytes, indicating start and finish of message, handling lost or corrupted data.

• A Data Link Layer for reliable transmission between endpoints

Problem

Communicate between any pair of a mesh of "nodes"



• Problems: find paths ("routes") between nodes, recover from broken paths.

• A *Network Layer* to send between network nodes

Problem

• Guarantee reliable end-to-end transmission through network



• Problems: handle lost messages, divide large messsages into small ones.

Message Encapsulation

- As messages go down the protocol stack, each received Service Data Unit (SDU) has address and control information added, to become the outgoing Protocol Data Unit (PDU) (and vice versa when going up the stack).
- Some messages may be segmented, one SDU into several PDUs.



Data Communications Transfer Rates

- Data sent over a communication link is converted into a byte stream.
- How long does it takes to transmit data or, what is the data rate as seen by the user?
- A fast, high-speed data link is no guarantee that users will see data transferred at that rate, or at anything approaching it.
 - Students are often confused by the word "rate".
 - Rate can be defined (my dictionary) as "a ratio between two measurements".
 - Here it usually means how often something happens. "A data rate of 1 megabits per second" means that 1 million bits are handled or delivered in each second.
 - And most importantly, "1 megabyte" in a computer usually means 1,048,576 bytes, while "1 megabit" in data communications always means exactly 1,000,000 bits.

Bits, Bytes, and Octets

- Many communications standards and documents call the 8-bit data unit an "octet", rather than a "byte". There are two other matters which may cause confusion.
- As octets are usually transmitted least-significant bit first, documents often write the bits in transmission order, with the least significant bit on the left and most significant bit on the right – the reverse of the usual numeric conventions on computers.
- The bits are often numbered 1 8 rather than "computerese" 0 – 7.



Data Transmission

- 1. The sending computer prepares a block of bytes, of an appropriate length and delivers the block to the data communications software.
- 2. The communications software divides the user data into sizes more appropriate to communications.
 - Or, many small pieces of user data may assembled into one longer entity.
- 3. The "packet" from step 2 is placed in an "envelope" for transmission, with a header including addresses, message type, length and control, and a trailer with check information.
 - In a multi-layer protocol, which is usual, steps 2 and 3 may be repeated several times, leaving the user data buried deeply in a nest of headers and trailers. At each stage the packet has a header, and trailer, enclosing a message body.
- 4. The final packet, including its header and trailer are sent over the data link.

Data Transmission (cont.)

- 5. The receiver checks the whole packet for validity and correctness, and extracts the message body.
 - The body may be assembled into part of a larger message or broken into smaller ones as needed.
 - The resultant message or messages may be sent on to the final user or passed to another communications layer for further processing.
- 6. A reply or acknowledgement may be prepared to send back to the original sender.
 - This is just another communications message and may be combined with others on the return path.
 - It has all of the encapsulation and transmission times given above.

Data Transmission (cont.)

- Each one of these steps takes a certain amount of time.
- These times must be combined to obtain the total effective time to process the message.
- Some texts develop complicated formulæ to describe these times, but it is better to show how to derive them in the general case, and then apply the rules to a particular example.

Automatic Repeat Request (ARQ)

This is the simplest transmission protocol.

• The sender sends one message and waits for a reply, either "correct" or "error – send again", before sending the next message.

Assume that

- a user message of U bytes (octets) is enclosed in an envelope of (H+T) octets for sending over a data link at R bits/second.
- the reply (positive or negative acknowledgement) is A octets, again enclosed in an envelope of (H+T) octets.
- the distance between sender and receiver is *D*, and information travels at a velocity *V*.

Neglecting any processing delays, we can calculate

- The time to get a reply, and
- The time before the next message may be sent.

 $\frac{8(H+U+T)}{R}$ Time to send packet from sender = Time for packet to travel to receiver = Time for receiver to send reply = $\frac{8(H + A + T)}{R}$ Time for reply to return to sender = $\frac{\sim}{V}$ $\frac{8(2H+A+U+2T)}{R} + \frac{2D}{V}$ Total time before next message may be sent = Effective or visible user data rate to send U bytes = $\frac{8(\overline{2H+A+U+2T})}{R} + \frac{2D}{V}$ $\frac{8(2H+A+U+2T)}{2}$ Message generation time = R Link delay or link latency =

There are two components here.

- The message generation time is proportional to the length of the total data transmitted (= user + ack), including overheads, and is inversely proportional to the bit-rate *R* on the link.
 - For long messages, the transmission time (U+A)/R may be reduced significantly by choosing a higher link bit rate R.
 - For very short user messages, the overheads 2(*H*+*T*) may become important.
- Choosing a "faster link" affects only the message generation time, not the link delay or latency.
 - The data can be put into the link at more bits per second if *R* is increased.
- A faster link NEVER affects the transmission velocity or link latency.
 - The speed of light is a fundamental constraint.

- The transmission delay or link latency is dependent only on the distance or link length *D*.
- The transmission velocity V is almost always 300,000 km/s = 3×10⁸ m/s = speed of light c, for radio 200,000 km/s = 2×10⁸ m/s = 2/3c, for cables (either copper or optical)
- For transoceanic links the latency can be a very important value indeed!

Example

Assume:

Reply message is 30 octets plus envelope = 240 bits User data size U = 1000 bytes = 8000 bits Header/trailer overhead (T+H) = 30 octet = 240 bits End to end cable length D = 1000 metres Signalling data rate R = 10 Mbit/s = 10^7 bit/s Signal velocity in cable V = 2×10^8 m/s

Then

End-to-end latency = $5 \mu s$

Time to send message = $8 \times 1030/10^7$ = $824 \mu s$ Time to send reply = $8 \times 60/10^7$ = $48.0 \mu s$

Total time to send 1000 user bytes: (send + outward + reply + reverse) = 882.0 μs

Effective user data rate = 1.134 byte/µs

Compare with a naive prediction:

Ignoring protocol & latency, the user data rate is $10^{7}/8 = 1.25$ byte/µs Protocol and latency cause a 9.3% reduction in performance.

- This example is for a relatively short and (by modern standards) slow link with a moderately large packet.
- A longer link, faster signalling rate and a smaller packet would give a much larger reduction in performance compared with the raw link speed (the "naïve prediction"), for ARQ protocols.
- Extreme precision is seldom needed, or even appropriate, in our examples. Packet or data sizes vary widely and cable velocities are seldom known to better than ±1% anyway.
- What is important is that you learn how to form (and justify) a good estimate of the effects of latency, packet overheads, etc.

Systems with overlapped transmission

- Send-and-wait (*e.g.* ARQ) signalling systems or protocols are unsuitable for high speed transmission with high latency (= long distance).
- It is better to send packets without waiting for a response, unless there are a lot of "lost packets".
- Replies will eventually come back to confirm that each packet has been received.
 - Each packet must be buffered at the sender, until it is acknowledged!
- The data rate is now close to R×U/(U+T+H), but we need to have big enough buffers. (How big?)

Example

 A transoceanic fibre-optic cable (D = 10,000 km) transmits ATM cells (U = 48, H = 5, T = 0) at 620 Mb/s.

- How much user data must retained at the sender?

- The cable velocity is 200,000 km/s;
 - with a round-trip distance of 20,000 km, the round-trip delay is 0.1 s.
- The largest possible user data rate on the cable (if packets are being sent continuously) is

$$\frac{620}{8} \times \frac{48}{53}$$
 octet/µs = 70.2 octet /µs

- In 0.1 s, we send 7.02×10^6 octets.
- The buffer at the sender should be at least 7 MB (= 7 × 1,048,576 bytes), so that any lost packets can be resent in case of error.

Consider the Network Layer

- The sending Network Layer receives data from its Transport Layer and passes it down to its Data Link Layer for transmission.
- The receiving Network Layer receives data from its Data Link Layer and passes it up to its Transport Layer.
- The Network Layer is written to exchange messages with the other Network Layer (its "peer" layer), without worrying how that message exchange occurs.
- Provided that it presents a standard interface to the Transport Layer (above) and to the DataLink Layer (below) and performs its defined functions,
 - the Network Layer is independent of the layers above and below, and even of its peer layer.

Message Encapsulation (review)

- As messages go down the protocol stack, each received Service Data Unit (SDU) has address and control information added, to become the outgoing Protocol Data Unit (PDU) (and vice versa when going up the stack).
- Some messages may be segmented, one SDU into several PDUs.



Inter-layer Interface

- A given layer (layer *n*) exchanges Service Data Units (SDUs) with the layer above it (layer *n*+1).
- Layer *n* exchanges *Protocol Data Units* (PDUs) with the layer below it (layer *n*–1).
- Layer *n* assumes that it is exchanging PDUs with its corresponding peer layer *n* at the other end.
- At the boundary between layers, the same message is regarded as a PDU by layer *n* and a SDU by layer *n*–1.

At the Network Layer, the data units and communication paths are —



Note that the same message is a

- Protocol Data Unit (PDU) above the layer boundary, and a
- Service Data Unit (SDU) below the boundary.

Inter-layer Interface (cont.)

- The layers can be implemented as subroutines, but this gives an inflexible, synchronous system with little or no overlap of user functions and data transfer.
- It is better to implement layers as independent tasks or processes, communicating with neighbours through queues.
 Layer n+1 task or process



Layer *n* task or process

Multiple users and access points or ports

- There are often several users above a level, to allow multiple users of a communication path.
- Communication is through "access points" or "ports", each identified by an address or identifier. The address must be added to the message in the lower level.



Special Data link and Physical Layers

- The OSI model was designed for point-to-point links, carrying only one type of traffic.
- For Local Area Networks (LANs) and many other newer networks, layers 1 and 2 are modified, by splitting the Data link Layer into two.

LLC Logical Link Control
MAC Medium Access Control
PHY Physical

- **LLC** Provides a standard interface to the MAC, including multiplexing
- MAC Framing/deframing data units, error checking, coordinates access to medium (physical protocol)
- **PHY** Dependent on network type (MAC and PHY operate as a unit)
- IEEE 802.3 defines an optional LLC within the Ethernet data field. This is used by SNAP to provide backwards compatibility with other MAC protocols, such as DECNET.

Connectionless and Connection-Oriented Services

- Modern communication almost always uses small bundles of data called "packets" (100 – 10,000 bytes).
 - Links can be shared between many competing users.
- Connectionless services send packets (or datagrams) into the network, with no prior notice.
 - Most LAN and Internet services (IP) are connectionless.
 - Each packet must have a complete end-point address.
- *Connection-oriented* services first establish a "Virtual Circuit" from the source to the destination end-points.
 - Packets have only the V-C number; tables in intermediate nodes maintain the connections for the virtual circuit.
 - The source and destination may then use the V-C much as though there was a physical link.