2010 PART IIA 3F5 COMPUTER AND NETWORK SYSTEMS
DR T W DRUMMOND

ENGINEERING TRIPOS PART IIA

Day Date April 2010 9 to 10.30

Module 3F5

COMPUTER AND NETWORK SYSTEMS CRIB

Answer not more than ***three*** *questions.*

All questions carry the same number of marks.

The ***approximate*** *percentage of marks allocated to each part of a question is indicated in the right margin.*

There are no attachments.

STATIONERY REQUIREMENTS
Single-sided script paper

SPECIAL REQUIREMENTS
Engineering Data Book
CUED approved calculator allowed

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

1 (a) What are the advantages and disadvantages of virtual memory? [20%]

Virtual memory allows each process to have its own address space and keeps the address spaces of different processes separate which reduces security risks and bugs. It also allows processes to use memory footprints which are larger than the pysical memory of the machine by storing some memory pages on disk. The main disadvantage is that memory access requires calculation of the physical address which can be slow if memory needs to be accessed to perform this computation.

(b) Explain how the main disadvantage can be reduced by using a Translation Lookaside Buffer (TLB). [10%]

TLBs are a cache for the computation of physical memory locations. This allows the physical location to be computed without an access to main memory which speeds up the process of reading or writing to main memory by a factor of two.

(c) Assuming a 32-bit virtual address, a 30-bit physical address, with 4kB pages, sketch the design of a 2-way set-associative TLB with a block size of 2 and a total of 32 blocks. [40%]

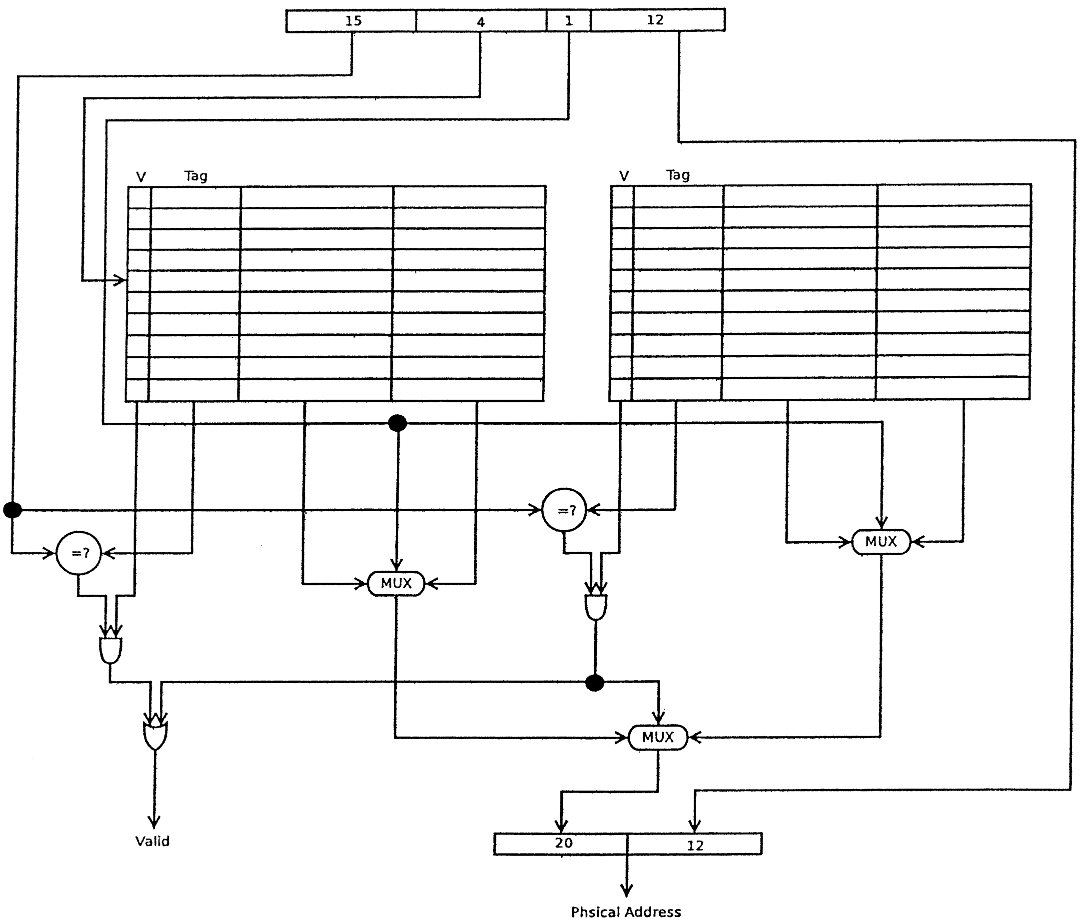

(d) Assume a MUX or testing equality takes two gate delays and that a table lookup takes six gate delays. If each gate causes a delay of 0.5ns, how long does your design of TLB take to compute a physical address?

For the design shown, this is 11 gate delays or 5.5ns.

Assuming that the memory cache has the same performance, how fast could the processor be clocked assuming a pipelined design? [20%]

Assuming the cache also takes 5.5 ns that makes 11 ns total for both, which would limit the clock speed to 909 MHz.

(e) What difference would be made if the cache indexed *virtual* addresses? What are the disadvantages of such a design? [10%]

If the cache indexed virtual memory, there is no need for the TLB when data is found in cache. This could potentially double the clock rate. However, this means that the cache must be flushed on every context change - or the cache has to be augmented with a process id to ensure correctness.

2 (a) What is meant by a data hazard? [10%]

A data hazard occurs when a register is needed in the CPU before it has been written to by a previous instruction.

(b) What is the difference between *caller* save and *callee* save? [10%]

In caller save, the calling function is responsible for putting registers on the stack to preserve their values. In callee save, the called function is responsible.

(c) The assembler code below shows a function which counts how many bits are set in an integer. For example, the integer 23 has the binary representation 10111 and hence if the value 23 were passed to `bitcount`, it would return the value 4.

```
# Function bitcount
# takes integer argument in register 4
# counts the number of bits set in this register
# and returns the result in register 2

bitcount:
    add $2,$0,$0          # store 0 in register 2 which will hold the result
    addi $20, $0, 32      # store 32 in register 20 (this is a loop counter)

loop:
    andi $21, $4, 1       # AND register 4 with 1 and store in register 21
                          # register 21 is now 1 if lowest bit of register 4 was set
    beq $21, $0, notset   # if the result is zero jump to notset:

    addi $2, $2, 1        # if we got here we found a bit set so add 1 to register 2

notset:
    srl $4, $4, 1         # shift register 4 right by one bit so we can test the next bit
    addi $20, $20, -1     # subtract 1 from register 20 (the loop counter)
    bgtz $20, loop        # if reg 20 has not reached zero yet, go round the loop again

    jr $31                # return from function (the result is in register 2)
```

Assuming no data forwarding, identify the data hazards in this code. Can any of these hazards be avoided by reordering the instructions? [30%]

The instructions `beq $21, $0, notset` and `bgtz $20, loop` both cause data hazards. They can both be avoided by moving the instruction `addi $20, $20, -1` to just before the beq instruction.

(d) By convention, registers 2 and 3 are used for the result of a function, registers 4–7 are used for the arguments of a function, registers 8–15 are caller save and registers 16–23 are callee save.

Following this convention, which registers would need to be saved by the function bitcount? Would it be possible to use different registers to avoid this? [20%]

Registers 20 and 21 are callee save - so these would need to be saved. Registers 8 and 9 could be used instead. These are caller save and since this function does not call any other functions, they are safe to use.

(cont.

(e) Write a function in MIPS assembly code which takes two integers as arguments (in registers 4 and 5), counts how many bits are set in both of them and returns the sum of those numbers in register 2. Your function should use the bitcount function already given. Minor errors in syntax will not be penalised, but you should clearly state what each instruction does. [30%]

```
bitcount64:
  add $16, $5, $0        # copy register 5 into register 16
  jal bitcount           # call the bitcount function to add up the bits in register 4
  add $17, $2, $0        # copy the result into register 17
  add $4, $16, $0        # copy register 16 (previously register 5) into register 4
  jal bitcount           # call the bitcount function to add up the bits in the other register
  add $2, $2, $17        # add on the other result
  jr $31                 # return from function
```

In order to be correct this function should store registers 16 and 17 on the stack.

3 (a) What are the three main features of the synchronous digital hierarchy (SDH) which helped to correct the defects in the older pleisiochronous digital hierarchy (PDH) used in modern telephone networks. Explain why these three features make SDH a far superior telephone network structure. [25%]

(b) Describe how an ideal SDH wide area network would be constructed across a country such as the United Kingdom and describe each component used to create such a wide area network.

Why is such a wide area network topology not possible in reality? [30%]

(c) The evolution of SDH also gave us integrated services digital network (ISDN) protocols and in particular the asynchronous transfer mode (ATM) protocol. Explain the basic structure of ATM and show how it was made compatible with the SDH system. [25%]

(d) Why is this SDH network now being replaced by internet using the voice over internet protocol set? [20%]

(TURN OVER

4 (a) Describe in detail the main differences between a circuit switch and a packet switch in terms of the requirements of a telecommunications network and a computer data network. [25%]

(b) One of the first packet switched protocols to be popular in the 1980s was X.25. Explain why the X.25 packet switched network protocol was successful and why it eventually evolved into frame relay. Identify the key features which were retained in the new frame relay protocol. [30%]

(c) One of the most important local area network (LAN) protocols in today's computer networks is the 802.3 standard, ethernet. Describe the original features of the 802.3 ethernet protocol for a coaxial cable based LAN. Why did this protocol evolve into the protocol that is in use today? [30%]

(d) Given how the protocols have evolved in parts b) and c) above, compare the relevant features of modern day frame relay and ethernet. [15%]

END OF PAPER

Q1.

a) What are the three main features of the synchronous digital hierarchy (SDH) which helped to correct the defects in the older pleisiochronous digital hierarchy (PDH) used in modern telephone networks. Explain why these three features make SDH a far superior telephone network structure. [25%]

b) Describe how an ideal SDH wide area network would be constructed across a country such as the United Kingdom and describe each component used to create such a wide area network.

Why is such a wide area network topology not possible in reality? [30%]

c) The evolution of SDH also gave us integrated services digital network (ISDN) protocols and in particular the asynchronous transfer mode (ATM) protocol. Explain the basic structure of ATM and show how it was made compatible with the SDH system. [25%]

d) Why is this SDH network now being replaced by internet using the voice over internet protocol set? [20%]

Q1 Crib – more verbose than expected

a) The aim of SDH was to correct the defects of PDH.

- Synchronisation. All elements of the system are synchronised to the same master clock. All tributaries are at a common rate before Mux. Hence only one level of bit stuffing is required.
- Pointers. Information about the muxed signals is transmitted at fixed intervals, which indicate the positions of units within the mux process. Hence any unit within the mux process can be identified and drop and insert processes can be done dynamically from an external viewpoint.
- Control and management. This is done outside the transmission process allowing complete control of the mux process. Time slots are put aside for a variety of tasks in ensuring the synchronicity of voice and data services.

We can now provide synchronisation and network management to the SDH network, which is entirely software controllable. Networks can be invisibly configure if there is a failure. Resources can be dynamically allocated to allow maintenance. Network capacity can be dynamically allocated to allow high bandwidth users like video conferencing or LAN access. Room for expansion and future-proofing.

b) An idealised synchronous transmission network contains not only SDH fibre optic rings, but also different protocol systems like submarine cables, satellite and microwave.

In this system, the synchronous OXC is used to connect together a series of self healing rings which also contain add/drop Mux units.

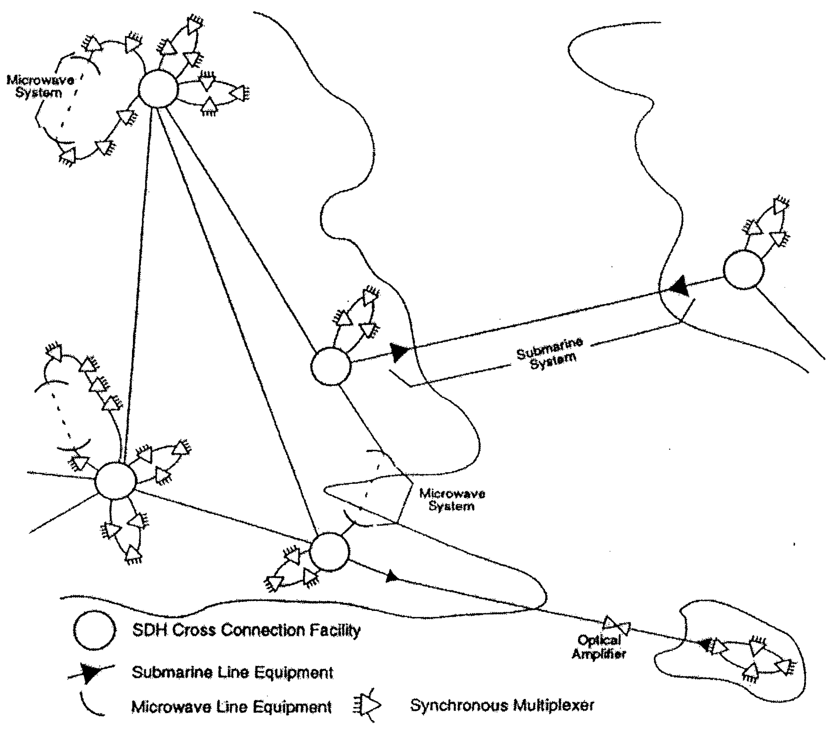

Synchronous multiplexer.

Allows optical integration at the STM-N level. Optical interfaces are often repeated for redundancy as main/standby or East/West pairs. Often in a ring topology. Must offer digital mux to STM-1 with flexible inputs.

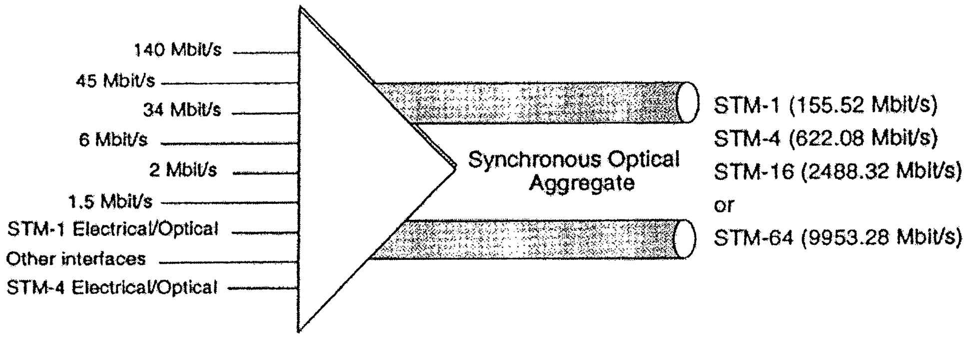

Synchronous, Plesiochronous, LANs/FDDI and MAN access, B-ISDN and ISDN, ATM, IP

The synchronous cross-connect (TDM switch)

These are used to reconfigure the network by setting up semi-permanent connections at the VC level. This is mostly a time slot re-arrangement operation, as the VC's are 'virtual' containers. It could also involve space switching between different network nodes. Connections are set by the network manager, not by the user.

This not the case for a real SDH network as there was an existing PDH system which had to be replaced gradually. The PDH system was built on existing local office exchange systems, hence the SDH is far from ideal and has a residual tiered structure due to the older systems.

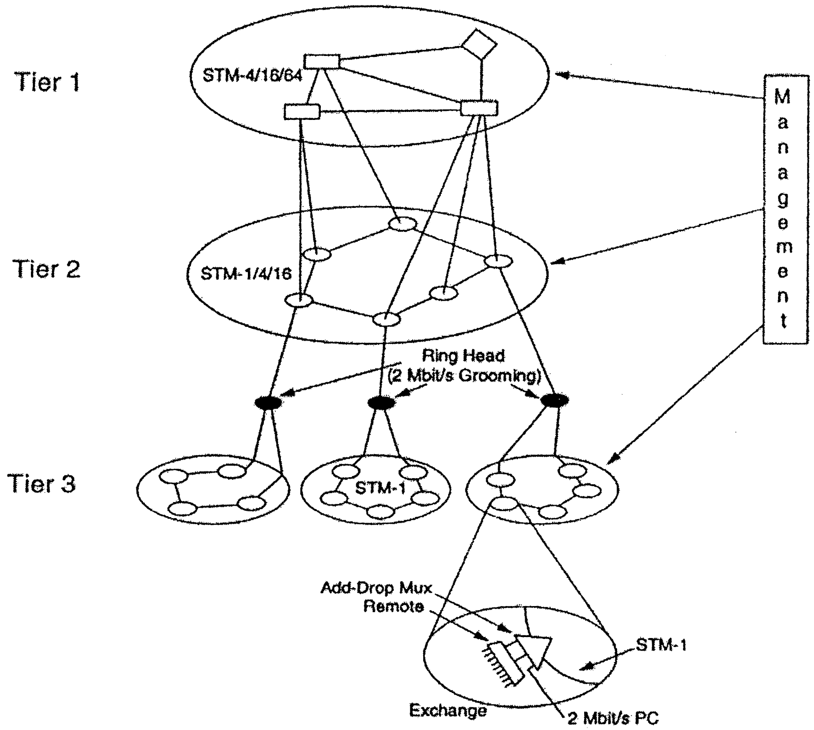

One solution is to replace the local exchanges with small remote concentrators (mux units) and connect these to the main exchanges with synchronous links. Local exchanges in tier 1 are replaced with an add/drop mux and remote concentrators working up to the STM-1 level. The mux units are used to form a basic ring structure.

c) Along with the development of SDH came the question; how can we accommodate both voice and data on SDH? SDH has been designed for low latency and minimal delays, which is ideal for voice. Data tends to be bursty and has been suited to packet switched networks. This more difficult to map onto SDH, especially compressed video streams which are bursty and delay sensitive.

For this reason, the Integrated Services Digital Network (ISDN) standard was defined to provide digital access to the home. This is the integration of both telephony and data into the same network. But this network will eventually have to cope with extreme data rates and bursty systems. We need a network which is capable of emulating both packet switching and circuit switching. We need to emulate both packet and circuit performance to provide both bandwidth on demand as well as constant bandwidth services.

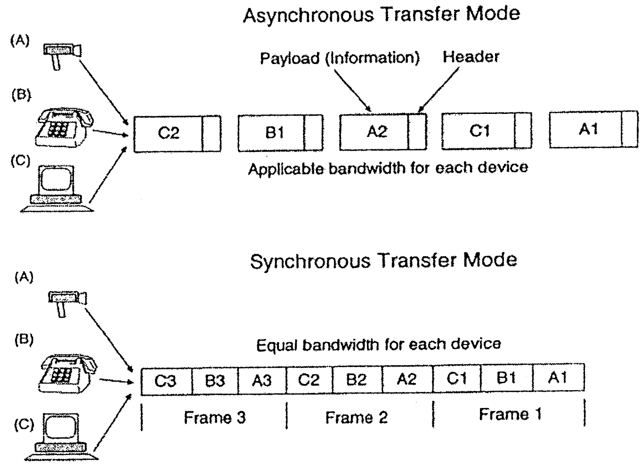

One possible solution is B-ISDN called asynchronous transfer mode or ATM.

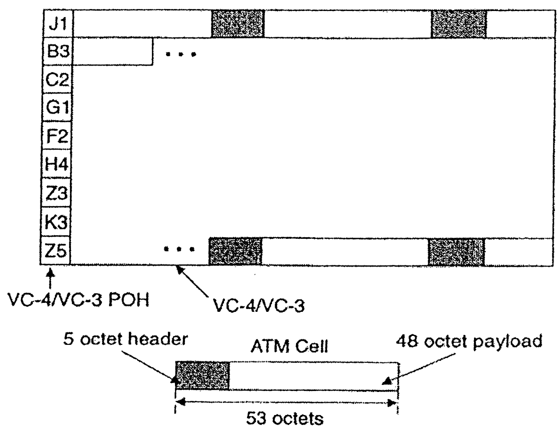

The advantage of the STM-1 frame is that it offers a huge wealth of virtual bandwidth as well as management services which will make ATM or IP over SDH easier than just pure ATM or IP systems. In ATM, the cell is limited to 53bytes, a 48 byte payload and a 5 byte header. ATM is defined at a transmission rate of 155.52Mbits/sec to be compliant with SDH and SONET.

d) There are many correct answers to this part of the question. Voice services are one of the best models for charging for network BW. Plus they can be expanded to multimedia services, then there is an even better cost model. Modern telcos want to offer multiple services to customers to expand their service base. Essentially VoIP allows a competitive edge over SDH in the commercial realm. The BW overhead of voice signal is so small that the internet can offer a virtually connection oriented service (delay free), hence it is possible to run a delay minimised VoIP service. The SDH system is near perfect for voice services, but it is controlled by a small number of companies, hence VoIP has proliferated as an alternative to the monopoly of the SDH providers.

Q2.

a) Describe in detail the main differences between a circuit switch and a packet switch in terms of the requirements of a telecommunications network and a computer data network.[25%]

b) One of the first packet switched protocols to be popular in the 1980's was X.25. Explain why the X.25 packet switched network protocol was successful and why it eventually evolved into frame relay. Identify the key features which were retained in the new frame relay protocol. [30%]

c) One of the most important local area network (LAN) protocols in today's computer networks is the 802.3 standard, ethernet. Describe the original features of the 802.3 ethernet protocol for a coaxial cable based LAN. Why did this protocol evolve into the protocol that is in use today? [30%]

d) Given how the protocols have evolved in parts b) an c) above, compare the relevant features of modern day frame relay and ethernet. [15%]

Q1 Crib

a) A circuit switch is normally used in telecommunications networks, where minimising delay across the network are important. This is the case when considering either video or voice services within a network. A circuit switch is a connection oriented switch which means that the connection is guaranteed for the duration of the call procedure. In the case of a voice service, this means that the connection is maintained throughout the duration of the call. This originally meant that there was an individual 'circuit' used for each possible connection, however in the modern SDH system this means that there is a guaranteed 64kbits/sec channel held open for the duration of the call. One of the key features of a circuit switch is that it is always controlled by the external management or control system of the network and does not depend on the data that is passing through it. If computer data is passed through a circuit switch, then there is a large amount of wasted bandwidth.

A packet switch is a switch which is set by the data passing through it and it can be either connection oriented (as is the case when virtual channels are allocated) or connectionless as is the case when datagrams are used (as in the internet). A packet switch usually controls access on and off a high bitrate medium via statistical multiplexing using either predefined virtual or logical channel, or by using the routing information contained within each packet. A key point in packet switches is that the switch state depends on the data passing through it, hence if there is a bottleneck, then data must be buffered and therefore it is very difficult to control packet delays. For computer date this is not a problem, however for voice and video service, this is a serious limitation.

b) X.25 was one of the first data protocols to be well defined and OSI compliant, hence it forms the basis from which many later transport protocols have been developed. A key feature was the frame check sequence (FCS) as old co-axial cable lines were very poor quality and very prone to errors. Today's digital transmission is several orders of magnitude better quality than that of the 1970s, so that the heavy duty error detection and correction techniques used by X.25 have become redundant. Windowing and acknowledgement are now largely superfluous. Instead of it being undertaken by the network, the job of error control or recovery is left to higher layer protocols.

The frame format consists of five basic information fields, much like the data link layer format of X.25. The flag marks the beginning of the frame, delineating it from the previous frame.

The address field carries the data link connection identifier (DLCI), the equivalent of the logical channel number (LCN) (ie is an OSI layer 2 address).
The control field contains supervision indication for the connection, like receiver ready (RR) and not ready (RNR), etc. It also contains the length of user information.
The information field contains user information. This may be up to 65536 bytes in length.
Finally the **frame check sequence** (FCS) is a cyclic redundancy check (CRC) code providing error detection.

c) CSMA/CD (Ethernet) is a contention based protocol for use with a shared media such as coaxial cable. On a CSMA/CD LAN the terminals do not request permission from a central controller before transmitting data onto the transmission channel; they contend for its use.

Before transmitting a packet of data, a sending terminal 'listens' to check whether a path is in use and if so, it waits before transmitting its data. Even when it starts to send data, it needs to continue checking the path to make sure that no other stations have started sending data at the same time. If the sending terminal's output does not match that which it is simultaneously monitoring on the transmission path, it knows there has been a collision. To receive data, the medium access control (MAC) software in each terminal monitors the transmission path, decoding the destination address of each packet passing through to find out whether it is the intended destination. If it is, the data is decoded, if not the data is ignored. Theory suggests that random collision of a large number of devices can lead to transmission degradation under heavy traffic.

LAN hubs (or switches) have become necessary as the data rate is too high for shared media. This caused the development of LANs with structured cabling, using hubs and twisted pair (10 and 100baseT), in a star configuration.

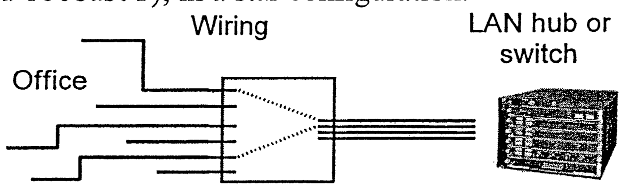

When installed as part of a structured cabling scheme (nowadays the most common realisation of ethernet), twisted pair or coaxial cabling provides for the transmission medium. Multiple twisted pair cables are usually installed in each individual office and near each desk, and are wired back to a wiring cabinet. Next to the wiring cabinet is a LAN hub which replaces the coaxial backbone, so the arrangement is often called a collapsed backbone. The bus topology still exists, but only within the hub itself. If new devices are added, then it is patched through the wiring cabinet.

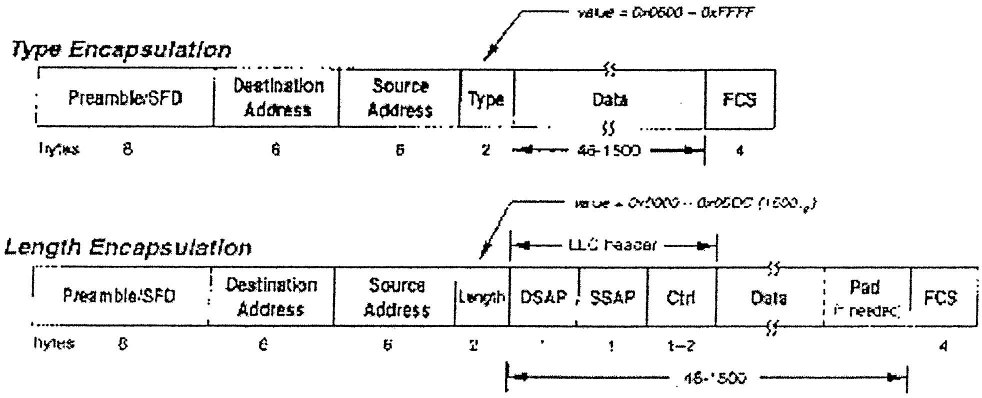

After removing the collision mechanism we are basically left with a transport frame format.

d) The two protocols have evolved into the same basic structure. Using the frame relay format as an analogy we have:

The address field is equivalent to the ethernet MAC address.
The control field is equivalent to the ethernet type field, especially in the LLC version of ethernet (length encapsulation) with SAPs.
The information field is common to both protocols, however there is one important difference. Frame relay MTU = 64k bytes, ethernet MTU = 1.5k bytes.
Frame check sequence (FCS) is very similar in both frame formats.