(1) Secondary Storage:-

Large storage requirements of most computer systems are economically realized in the form of magnetic disks, optical disks and magnetic tapes which are usually refferred to as secondary storage devices.

i) Magnetic Hard Disks:-

- The storage medium in a magnetic-disk system consists of one or more disks mounted on a common spindle. A thin magnetic film is deposited on each disk, usually on both sides.
- The disks are placed in a rotary drive so that the magnetized surface move in close proximity to read/write heads.
- Digital information can be stored on the magnetic film by applying current pulses of suitable polarity to the magnetizing coil.
- The modern approach is to combine the clocking information with the data. Several different techniques have been developed for such encoding. One simple scheme, depicted is known as phase encoding or Manchester encoding.
- The drawback of Manchester encoding is its poor bit-storage density.
- The space required to represent each bit must be large enough to accommodate two changes in magnetization.
- Read/write heads must be ~~mataint~~ maintained at a very small distance from the moving disk surfaces in order to achieve high bit densities and reliable read/write operations.
- When the disks are moving at their steady rate, air pressure develops between the disk surface and the head and forces the head away from the surface.
- The above force can be counteracted by a spring-

loaded mounting arrangement for the head that allows it to be pressed toward the surface.

- The flexible spring connection between the head and its arm mounting permits the head to fly at the desired distance away from the surface in spite of any small variations in the flatness of surface.

- In most modern disk units, the disks and the read/write heads are placed in a sealed, air-filtered enclosure, this approach is known as Winchester technology.

- 2nd advantage of Winchester technology is that data integrity tends to be greater in sealed units where the storage medium is not exposed to contaminating elements.

- The disk system consists of three key parts.
  1st part is the assembly of disk platters, called disk.
  2nd part is electromechanical mechanisms that spins the disk and moves the read/write heads, called disk drive.
  3rd part is the electronic circuitry that controls the operational of the system, called disk controller.

Organization and Accesing of Data on a Disk

- The surface is divided into concentric tracks, and each track is divided into sectors.

- The set of corresponding tracks on all surfaces of a stack of disks forms a logical cylinder.

- The data are accessed by specifying the surface number, the track number, and the sector number.

- The read and write operations start as sector boundaries.

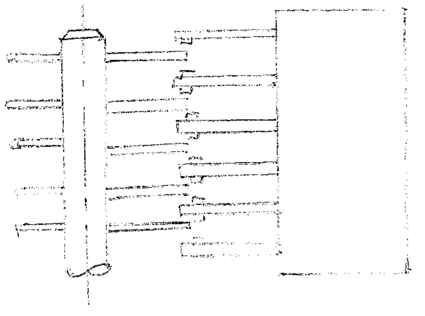

a) Mechanical Structure

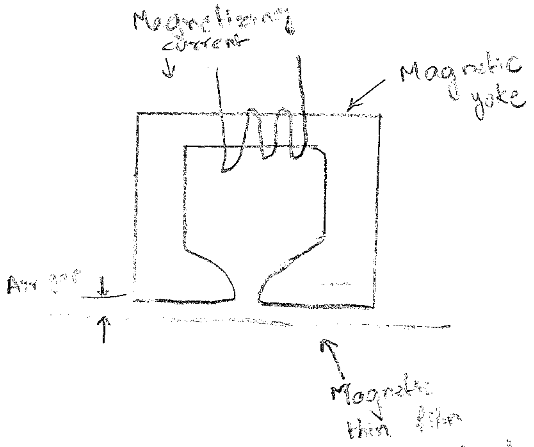

b) Read/write head detail

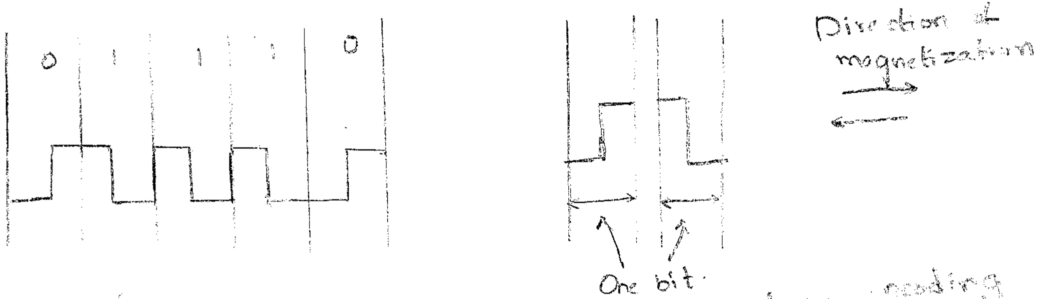

c) Bit representation by phase encoding

Magnetic Disk Principles

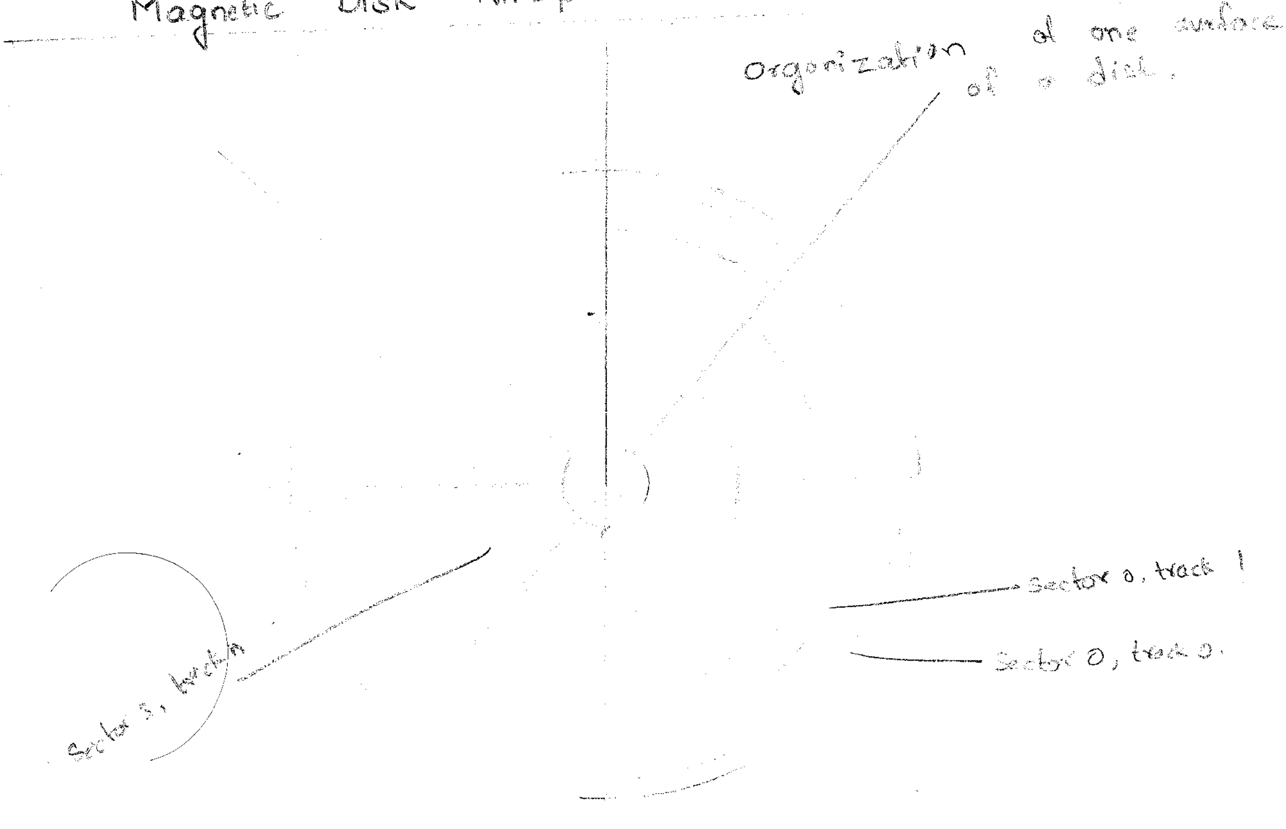

- Data bits are stored serially on each track. Each track usually contains 512 bytes of data, (other size be used).
- The data are preceded by a sector header that contains identification informatio (i.e. address) used to find the desired sector on the selected track.
- There are some additional bits that constitute the error correcting code(ECC), which are used to detect and correct errors that may have occured in writing or reading of the specified data types.
- To easily distinguish between two consecutive sectors, there is a small intersector gap.
- An unformatted disk has no information on its tracks.
- The formatting process divides the disk physically into tracks and sectors.

Access Time :-

- There are two components involved in the time delay between receiving an adress and the beginning of the actual data transfer. The first, called the seek time, is the time required to move the read/write head to the proper track.
- It depends on the initial position of the head relative to the track specified in the adress.
- The 2nd component is the rotation delay, also called latency time. (It is the amount of time that elapses after the head is positiond over the correct track until the starting position of the addressed sector passes under the read/write head.

Disk Controller :-

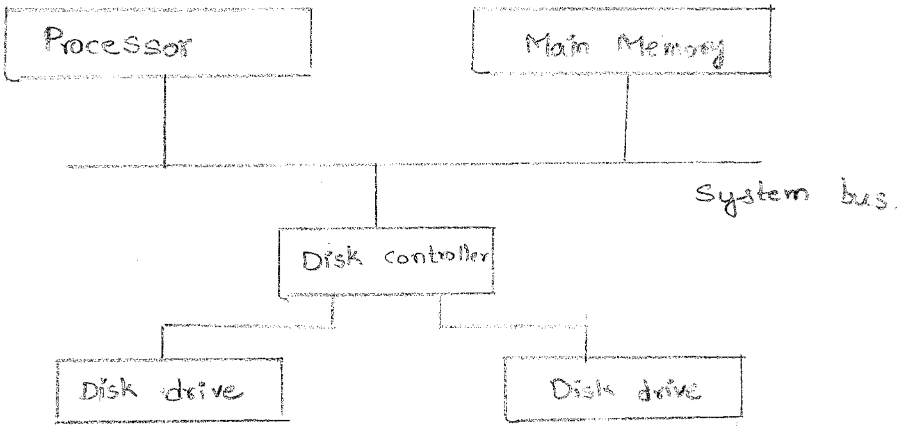

Disks connected to the system bus.

- The SCSI bus is capable of transferring data at much higher rates than the rate at which data can be read from disk tracks.
- The buffer is a semiconductor memory, capable of storing a few megabytes of data.
- Operation of a disk drive is controlled by a disk controller circuit, which also provides an interface between the disk drive and the bus that connects it to the rest of the computer system.
- The disk controller may be used to control more than one device.
- The OS initiates the transfers by issusing Read and Write requests, which entail loading the controller's registers with the necessary addressing and control information, typically, which entail loading the controller's registers with the necessary addressing and control information,

Main memory address - The address of the first main memory location of the block of words involved in the transfer.

Disk address - The location of the sector containing the beginning of the desired block of words.

Word count - The number of words in the block to be transferred.

⇒ On the disk drive side, the controller's major functions are:-

Seek - causes the disk drive to move the read/write head from its current position to the desired track.

Read - Initiates a Read operation, starting at the address specified in the disk address register. Data read serially from the disk are assembled into words and placed into the data buffer for transfer to the main memory. The number of words is determined by the word count register.

Write - Transfers data to the disk, using a control method similar to that for the Read operations.

Error checking - Computes the error correcting code (ECC) values for the data read from a given sector and compares it with the corresponding ECC value read from the disk.

Software and Operating System Implications:-

- All data transfers activities involving disks are initiated by the operating system. The disk is a non-volatile (memory) storage medium, so the OS itself is stored on a disk.

- During normal operation of a computer, parts of the OS are loaded into the main memory and executed as needed.

- Disk accesses are very slow compared to main memory access, mostly due to long seek times.
- After the OS initiates a disk transfer operation, it normally attempts to switch execution to some other task, to make use of the time it would otherwise spend waiting for the transfer to complete. The disk controller informs the OS when the transfer is completed by raising an interrupt.

Floppy Disks:-

- The Floppy disks are smaller, simpler and cheaper disk units that consist of a flexible, removable, plastic diskette coated with magnetic material.
- A hole in the center of the diskette allows a spindle mechanism in the disk drive to position and rotate the diskette.
- The simplest schemes used in the first floppy disks for recording data is phase or Manchester encoding mentioned earlier.
- A more complicated variant of this scheme, called double density is most often used in current standard floppy disks. It increases the storage density by a factor of 2 but it also requires more complex circuits in the disk controller.
- The main feature of floppy disks is their low cost and shipping convenience.
- Larger super-floppy disks are also available. One type of such disks known as the zip disk, can store more than 100 Mbytes

RAID Disk Arrays:-

- In 1988, researchers at the University of California - Berkeley proposed a storage system based on multiple disks [5]. They called it RAID, for Redundant Array of Inexpensive Disks.

- Six different configuration were proposed. They are known as RAID levels even though there is no hierarchy

~~Raid~~

- RAID 0 is the basic configuration intended to enhance performance. A single large file is stored in several seperate disk units by breaking the file up into a no. of smaller pieces and storing these pieces on different disks. This is called data striping.

- RAID 1 is intented to provide better reliability by storing identical copies of data on two disks rather than just one. The two disks are said to be mirrors of each other. If one disk drive fails, all read and write operations (take place) are directed to its mirror drive.

- RAID 2, RAID 3, RAID 4 levels achieve increased reliability through various parity checking schmes without requiring a full duplication of disks. All of the parity information is kept on one disk.

- RAID 5 also makes use of a parity-based error-recovery scheme. However, the parity information is distributed among all disks, rather than being stored on one disk.

- Dell Computer Corporation offers products based on RAID 0, 1, 5, 10.

Optical Disks:-

- Largest storage devices can also be implemented using optical means.
- The familiar compact disk (CD), used in audio systems, was the first practical application of optical ~~thee~~ technology.
- It is also used in computer environment to provide high-capacity read only memory (ROM).
- To provide high-quality sound recording and reproduction, 16-bit samples of the analog signals are taken at a rate of 44,100 samples per second.

CD Technology:-

- The optical technology is used for CD systems is based on a laser light source.
- A laser beam is directed onto the surface of the spinning disk. The physical identations in the surface are arranged along the tracks of the disk.
- A cross-section of a small portion of a CD is explained as follows. The bottom layer is polycarbonate plastic, which functions as a clear glass base. The surface of this plastic is programmed to store data by indenting it with pits. The unindented parts are called lands. A thin layer of reflecting aluminium material is placed on top of a programmed disk. The aluminum is then covered by a protective acrylic. Finally, the top most layer is deposited and stamped with a label. The total thickness of the disk is 1.2 mm. Almost all of it is contributed by the polycarbonate plastic. The other layers are very thin.
- The laser source and the photodetector are positioned below the polycarbonate plastic. The emitted beam travels through this plastic, reflects off the aluminum layer, and travels back toward the photodector. Note that from the laser side, the pits actually appear as bumps with respect to the lands.

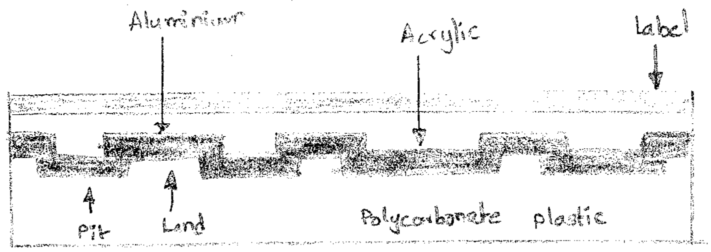

cross section

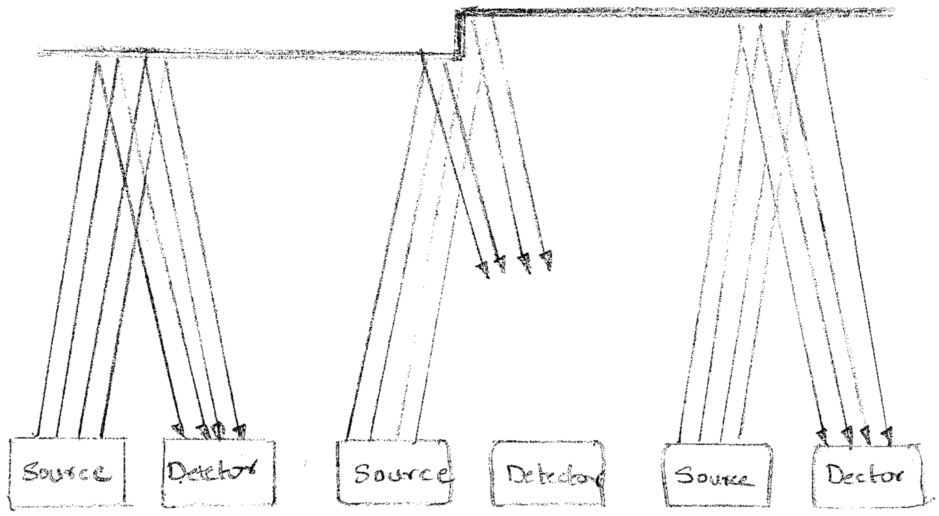

Transition from pit to land.

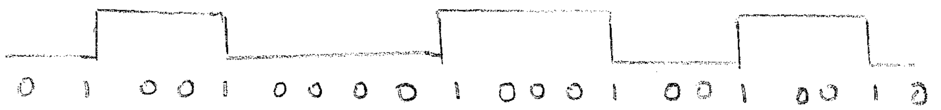

stored binary pattern

Optical disk.

CD-ROM :-

- Stored data are organized on CD-ROM tracks in the form of blocks that are called sector.
- There are several different formats for a sector. One format, known as Mode 1, uses 2352-bytes sectors. There is a 16-byte header that contains a synchronization field used to detect the beginning of the sector and addressing information used to identify the sector. This is followed by 2048 bytes of stored data. At the end of sector, there are 288 bytes used to implement the error-correcting scheme.
- Error ~~deection~~ detection and correction is done at more than one level.
- CD-ROM drives operate at a number of different rotational speeds. The basic speed, known as 1x, is 75 sectors per second.
- The importance of CD ROMs for computer systems stems from their large storage capacity and fast access times compared to other inexpensive portable media, such as floppy disks and magnetic tapes.
- Their attraction lies in the small physical size, low cost, and ease of handling as a removable and transportable mass-storage medium.

CD - Recordables :-

- A new type of CD was developed in the late 1990's on which data can be easily recorded by a computer user. It is known as CD-Recordable (CD-R). A spiral track is implemented on a ~~track~~ disk during the manufacturing process.
- A laser in a CD-R drive is used to burn pits into an organic dye on the track. When a burned spot is heated beyond a critical temp., it becomes opaque. Such burned spots reflect less light when subsequently read. The written data are stored permanently. Unused portions of a disk can be used to store additional data at a later time.

CD - Rewritables :-

- The most flexible CDs are those that can be written multiple times by the user. These are known as CD-RWs.

- The basic structure of CD-RWs is similar to the structure of CD-Rs. Intead of using an organic dye in the recording layer, an alloy of silver, indium, antimony and tellurium is used.

- The above alloy has interesting and useful behaviour when it is heated and cooled. If it is heated above its melting point (500 degrees C) and then cooled down, it goes into an amorphous state in which it absorbs light.

- But if it is heated only to about 200°c and this temp. is maintained for an extended period, a process known as annealing takes place, which leaves the alloy in a crystalline state that allows light to pass through it.

- The stored data can be erased using the annealing process, which returns the alloy to a uniform crystalline state.

- A reflective material is placed is placed above the recording layer to reflect the light when the disk is read.

- The CD-RW drive uses three different laser powers.

  i) The highest power is used to record the pits.

  ii) The middle power is used to put the alloy into its crystalline state, it is refferred to as the "erase power".

  iii) The lowest power is used to read the stored information

- But there is a limit on how many times a CD-RW disk can be rewritten (1000 times presently.

DVD Technology:-

- The physical size of a DVD disk is the same as for CDs. The disk is 1.2mm thick, and it is 120mm in diameter. Its storage capacity is made much larger than that of CDs by several design changes.
  - i) A red light laser with a wavelength of 635 nm is used instead of the infrared light laser used in CDs, which has a wavelength of 780 nm. The shorter wavelength makes it possible to focus the light to a smaller spot.
  - ii) Pits are smaller, having a minimum length of 0.4 microns.
  - iii) Tracks are placed closer together; the distance between tracks is 0.74 microns.

- Two single-sided disks can be put together to form a sandwich-like structure where the top disk is turned upside down. This can be done with single-layered disks, as specified in DVD-10, giving a composite disk with a capacity of 9.4 Gbytes. It can also be done with the double-layered disks, as specified in DVD-18, yielding a capacity of 17 Gbytes.

- Access times for DVD drives are similar to CD drives

- A rewritable version of DVD devices known as DVD-RAM, has also been developed.
- It has a large storage capacity but its only disadvantage is the higher price and relatively slow writing speed.

- To ensure that the data have been recorded correctly on the disk, a process known as write verification is performed by using DVD-RAM drive.

MAGNETIC TAPE SYSTEMS :-

- Magnetic tapes are used for off-line storage of large amounts of data.
- It uses the same principle as used in magnetic - disk recording.
- The main difference is that the magnetic film is deposited on a very thin 0.5 or 0.25 inch wide plastic tape.
- Data on the tape are organised in form of records seperated by gaps.
- Tape motion is stopped only when a record gap is underneath the read/write heads.
- The record gaps are long enough only ~~when a record gap~~ to allow the tape to attain its normal speed before the beginning of the next record is reached.

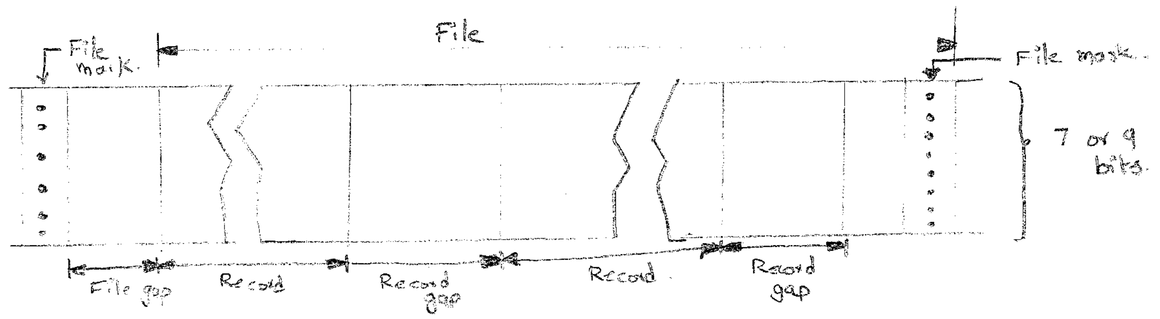

Organization of data on magnetic tape

- The controller of a magnetic tape drives enables the execution of a numbers of control commands in addition to read and write commands. Control Commands include following operations.

i) Rewind Tape.
ii) Rewind and unload tape
iii) Erase tape
iv) Write tape mark
v) Forward space one record
vi) Backspace one record.
vii) Forward space one file
viii) Backscape one file.

(2) Semi Conductor RAM Memories :-

- Semiconductor memories are available in a wide range of speeds.
- When first introduced in the late 1960s, they were much more expensive than the magnetic-core memories they replaced. Because of rapid advantages in VLSI (Very Large Scale Integration) technology, the cost of semiconductor memories has dropped dramatically.

Internal Organization of Memory chips :-

- Memory cells are usually organised in the form of an array, in which each cell is capable of storing one bit of information.
- Each row of cells constitutes a memory word, and all cells of a row are connected to a common line referred to as the word line, which is driven by the address decoder on the chip.
- The cell in each column are connected to a Sense/Write circuit by two bit lines.

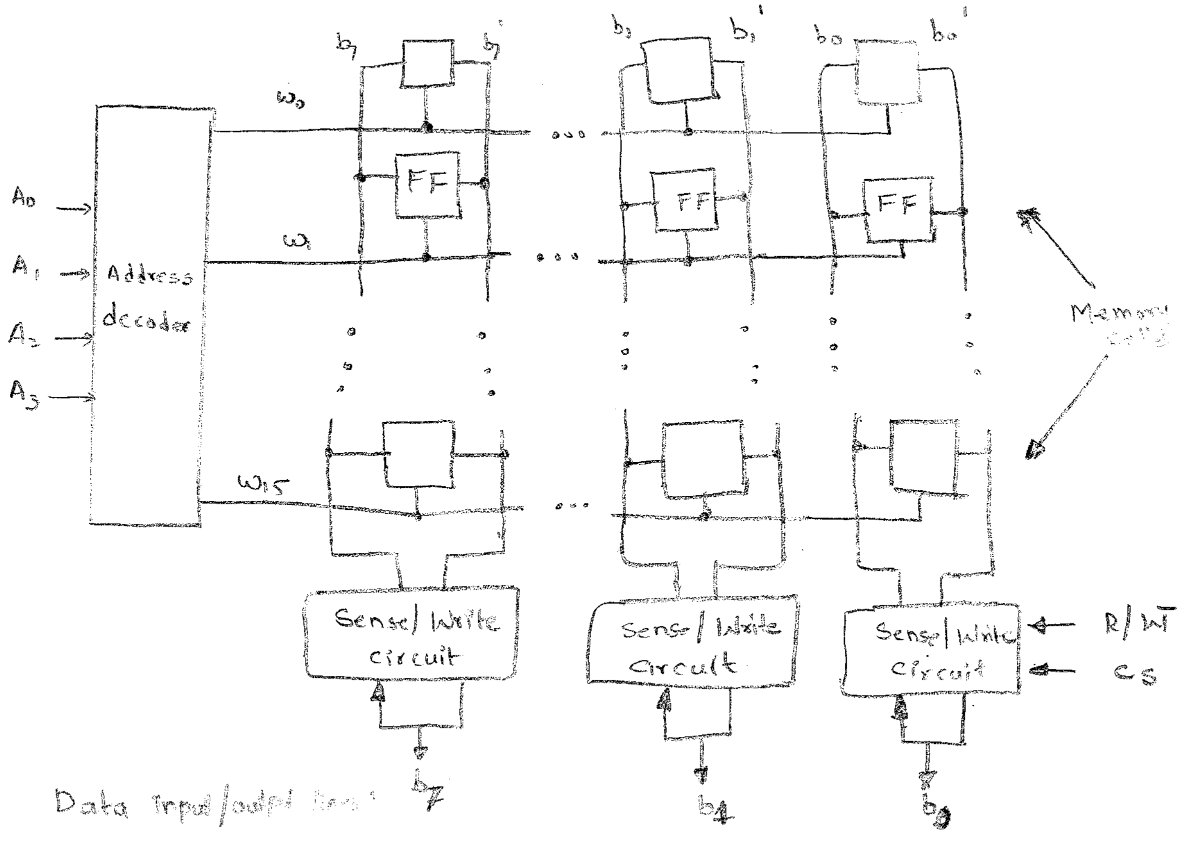

Organisation of bit cells in a memory chip

- The memory circuit stores 128 bits and requires 14 external connections for address, data, and control lines. Of course, it also needs two lines for power supply and ground connections.
- Consider now a slightly larger memory circuit, one that has 1K (1024) memory cells. This circuit can be organised as a 128 × 8 memory, requiring a total of 19 external connections.

STATIC MEMORIES :-

- Memories that consist of circuits capable of retaining their state as long as power is applied are known as static memories.
- Static RAM is implemented by two inverters are cross-connected to form a latch. The latch is connected to two bit lines by transistors T_1 and T_2. Those transistors act as switches that can be opened or closed under control of the word line.

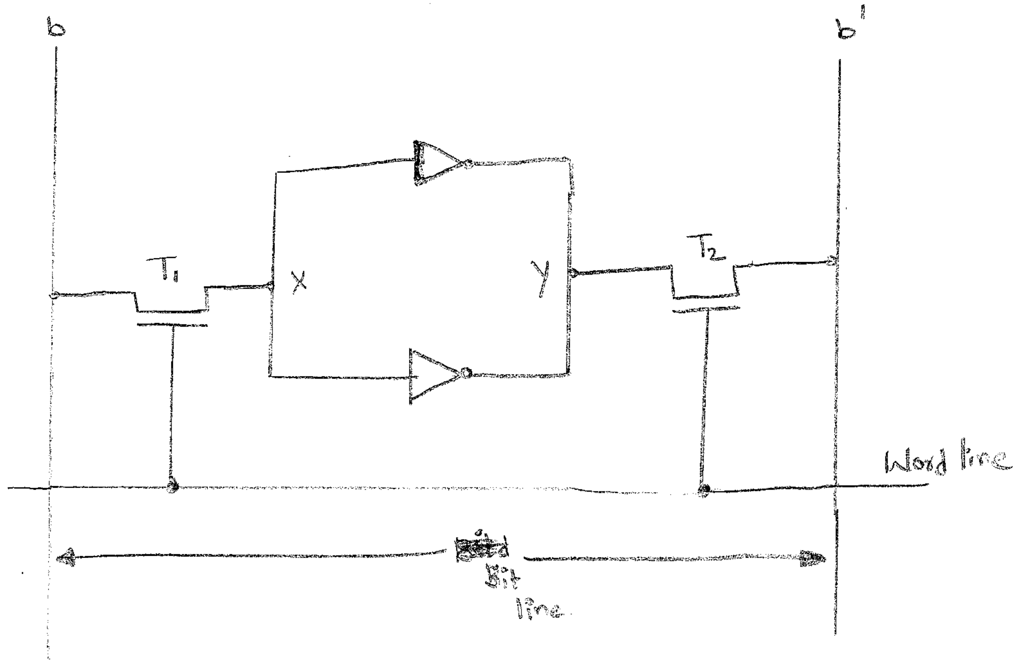

A SRAM cell

Read Operation

- In order to read the state of the SRAM cell, the word line is activated to close switches T_1 and T_2. If the cell is in state 1, the signal on bit line b is high and the signal on bit line b' is low. The opposite is true if the cell is in state 0. Thus b and b' are complements of each others. Sense/write circuits at the end of the bit lines monitor the state of b and b' and set the output accordingly.

Write Operation :-

The state of the cell is set by placing the appropriate value on bit line b and its complement on b', and then activating the word line. This forces the cell into the corresponding state. The required signals on the bit lines are generated by the sense/write circuit.

CMOS CELL :-

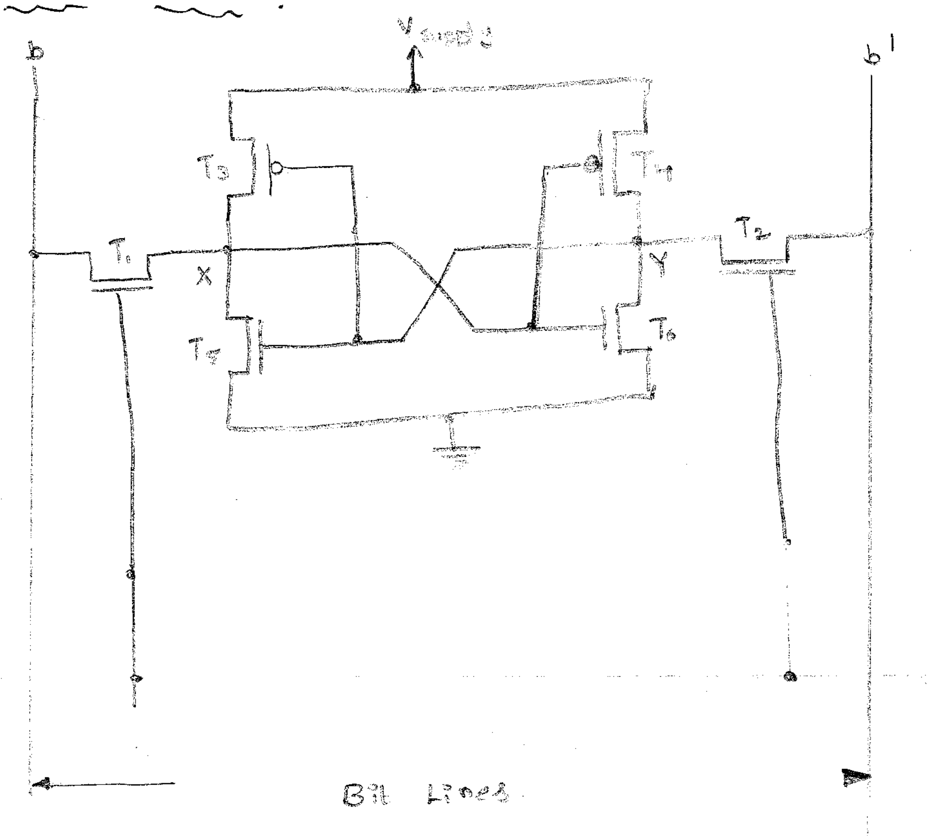

An example of a CMOS memory cell.

- A major advantage of CMOS SRAM is their very low power consumption because current flows in the cell only when the cell is being accessed. Otherwise, T_1, T_2 and one transistor in each inverter are turned off, ensuring that there is no active path between V_{supply} and ground.

Asynchronous DRAMS :-

- Information is stored in a dynamic memory cell in the form of a charge on a capacitor, and this charge can be maintained for only tens of milliseconds. Since the cell is required to store information for a much longer time, its contents must be periodically refreshed by restoring the capacitor charge to its full value.

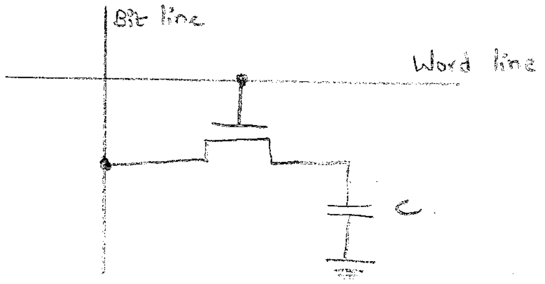

A single-transistor dynamic memory cell.

Fast Page Mode:-

- The most useful arrangement is to transfer the bytes in sequential order, which is achieved by applying a consecutive sequence of column addresses under the control of successive CAS signals. This scheme allows transferring a block of data at a much faster rate than can be achieved for transfers involving random addresses. The block transfer capability is reffered to as the fast page mode feature.

Synchronous DRAMS :-

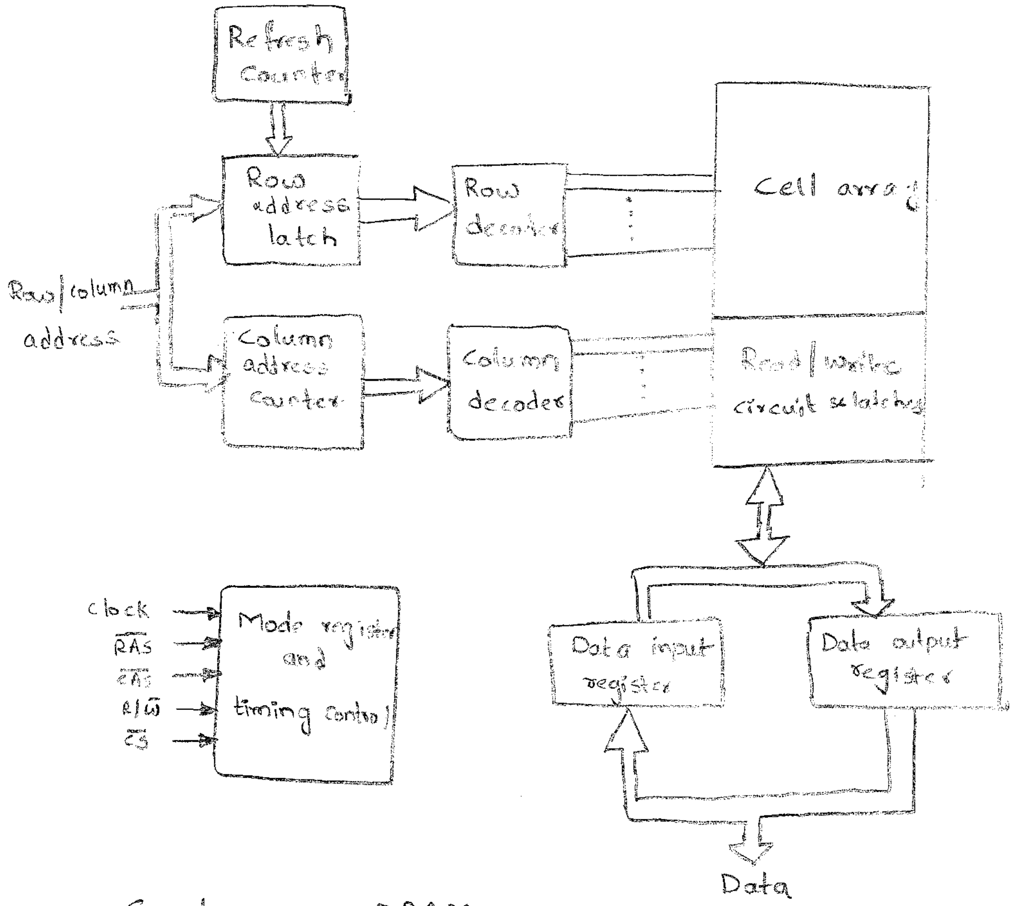

Synchronous DRAM

- The DRAMS whose operation is directly snchronized with a clock signal. Such memories are known as synchronous DRAMs.
- The cell array is the same as in asynchronous DRAMs.
- The address and data connections are buffered by means of registers.
- SDRAMS have several different modes of operations, which can be selected by writing control information into a mode register.
- SDRAMS have built in refresh circuitary. A part of this circuitary is a refresh counter which providies the addresses of the rows that are selected for refreshing.

- The term memory latency is used to refer to the amount of time it takes to transfer a word of data to or from the memory.
- When transferring blocks of data, it is of interest to know how much time is needed to transfer an entire block. Since block can be variable in size, it is useful to define a performance measure in terms of the number of bits or bytes that can be transferred in one second. This measure is often referred to as the memory bandwidth.

Structure of Larger Memories

1) Static Memory systems.
2) Dynamic Memory systems. (Memory modules usage).
   - Modern computers use very large memories. (small personal computer has at least 128M bytes of memory).
- A large memory leads to better performance because more of the programs and data used in processing can be held in the memory, thus reducing the frequency of accessing the information in Secondary storage.
- If a large memory is built by placing DRAM chips directly on the main system printed-circuit board that contains the processor, often referred to as a motherboard,
- Mother board will occupy an unacceptably large amount of space on the board.
- It is awkward to provide for future expansion of memory, because space must be allocated and wiring provided for the maximum expected size. These packaging considerations have led to the development of larger memory units known as SIMMs. and DIMMs. (Single In-line Memory Modules, Dual-In line Memory Modules).

Memory System Considerations:-

- Static RAMs are generally used only when very fast operation is the primary requirement.
- They are used mostly in cache memories.

Memory Controller:-

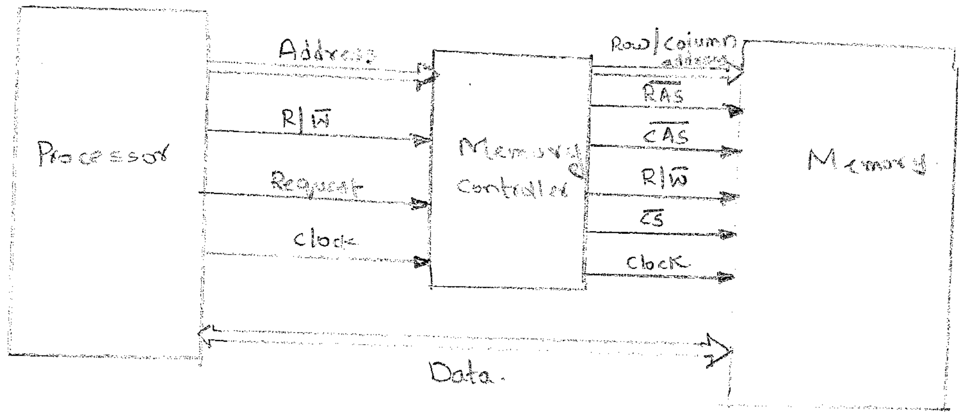

Use of a Memory Controller

- A typical processor issues all bits of an address at the same time. The required multiplexing of address bits is usually performed by a memory controller circuit which is interposed between the processor and the dynamic memory.
- The controller accepts a complete address and the R/$\overline{W}$ signal from the processor, under control of a Request signal which indicates that a memory access operation is needed.
- The controller then forwards the row and column portions of the address to the memory and generates the $\overline{RAS}$ and $\overline{CAS}$ signals
- It also sends the R/$\overline{W}$ and cs signals

Rambus Memory :-

- The performance of a dynamic memory is characterized by its latency and bandwidth.
- A very wide bus is expensive and requires a lot of space on a mother board.
- An alternative approach is to implement a narrow bus that is much faster. This approach was used by Rambus Inc. to develop a proprietary design known as Rambus.
- The key feature of Rambus Technology is a fast signaling method used to transfer information between chips.
- Rambus requires specially designed memory chips. These chips use cell arrays based on the standard DRAMS tech. Multiple banks of cell arrays are used to access more than one word at a time. Circuitary needed to interface to the Rambus channel is included on the chip. such chips are known as Rambus DRAMS
- RDRAM Chips can be assembled into larger modules, similar to SIMMs and DIMMs. One such module, called RIMM, can hold upto 16 RDRAMs.
- Rambus technology competes directly with the DDR SDRAM technology.
- A nontechnical consideration is that the specification of DDR SDRAM is an open standard, while RDRAM is a proprietary design of Rambus Inc. for which the manufactors of chips have to pay a royalty.

③ Performance :-

⇒ A common measure of success is the price/performance ratio.

- Performance depends on how fast machine instructions can be brought into the processor for execution and how fast (machine instructions ~~can~~ be brought into the) they can be executed.

i) Interleaving :-

- If the main memory of a computer is structured as a collection of physically seperate modules, each with its own address buffer register (ABR) and data buffer register (DBR), memory access operations may proceed in more than one module at the same. Thus the aggregate rate of transmission of words to and from the main memory system can be increased.

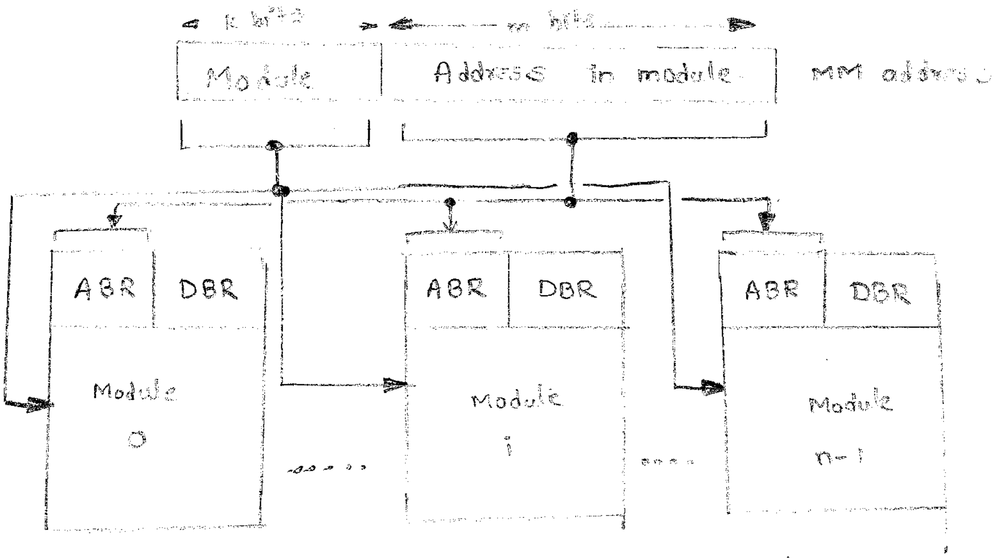

Consecutive words in a module

- The memory address generated by the processor is decoded in above figure.
- The high-order k bits name one of n modules, and the low order m bits name a particular word in that module. When consecutive locations are accessed, as happens when a block of data is transfered to a cache, only one module is involved.

- At the same time, however, devices with direct memory access (DMA) ability may be accessing information in other modules.

- The second and more effective way to address the modules in below figure.

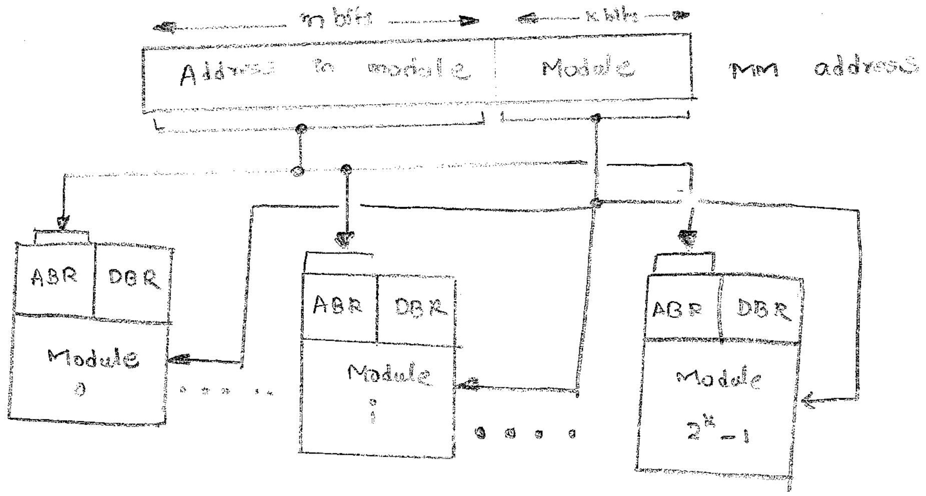

Consecutive words in consecutive modules.

- The above method is called memory interleaving.
- The lower-order k bits of the memory address select a module, and the high-order m bits name a location within that module.
- In this way, consecutive address are located in successive modes
- This results in both faster access to a block of data and higher average utilization of the memory system as a whole.
- To implement the interleaved structure, there must be 2^k modules. (else) otherwise there will be gaps of non existent locations in the memory address space.

ii) Hit Rate and Miss Penalty:-

- An excellent indicator of the effectiveness of a particular implementation of the memory hierarchy is the sucess rate in accessing information at various levels of the hierarchy.
- Recall that a successful access to data is called a hit.
- The number of hits states as a fraction of all attempted accesses is called the hit rate, and the miss rate is the number of misses stated as a fraction of attempted accesses.
- High hit rates, well over 0.9, are essential for high-performance computers.
- Performance is adversely affected by the actions that must be taken after a miss.
- The extra time needed to bring the desired information into the cache is called the miss penalty.
- In general, the miss penalty is the time needed to bring a block of data from a slower unit in the memory hierachy to a faster unit.
- This miss penalty can be reduced if the efficient mechanisms for transfering data between the various units of the hierachy are implemented.

iii) Caches on the Processor chip:-

- All high-performance processor chips include some form of a cache. Some manufactors have choosen to implement two seperate caches, one for instructions and another for data.

- A combined cache for instructions and data is likely to have a somewhat better hit rate because it offers greater flexibility in mapping new information into the cache.

- The disadvantages of seperate caches is that the increased parallelism comes at the expense of more complex circuitry.

- The average acess time experienced by the processor in a system with two levels of caches is

$$t_{ave} = h_1 C_1 + (1-h_1) h_2 C_2 + (1-h_1)(1-h_2) M$$

where

h_1 is the hit rate in the L1 cache.

h_2 is the hit rate in the L2 cache.

C_1 is the time to acees information in the L1 cache.

C_2 is the time to access information in the L2 cache.

M is the time to access information in the main memory.

- The number of misses in the L2 cache, given by the term $(1-h_1)(1-h_2)$, should be low. If both h_1 and h_2 are in the 90% range. Then the no. of misses will be less than 1% of the processor's memory accesses.