DATA RECORDING

Data are items of information. As we travel on the information superhighway, we find data all around us—weather data, sports data, stock data, sonar data from underwater objects, seismic data from an earthquake, satellite data from Mars, one-dimensional speech data, two-dimensional image data, three-dimensional video data, analog data, and digital data. Some data are natural, some are synthetic.

Data are recorded for later retrieval and use. Archeological findings are recorded to trace back the past. Current scientific findings are recorded to strive for the future. Musical data are recorded to listen at a later time on demand. So, what is data recording? Data recording refers to a system that can store data and restore it as faithfully as possible. A data recording device, therefore, involves both storage and retrieval mechanisms. A sample data recording system is shown in Fig. 1.

The source generates the raw data. Data acquisition is the process of measuring and gathering data from the source in some usable format. Acquired data is usually compressed to reduce the burden on the storage media. Each of these processes will be discussed next.

DATA ACQUISITION

As mentioned previously, data acquisition collects data in a usable format from the data source. This usually comprises a sensor, calibrator, or data converter. Figure 2 shows typical components of a data acquisition system. For example, for the recording of the temperature of a system, a thermocouple may

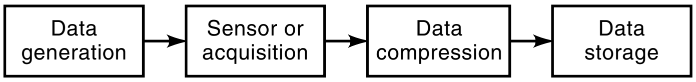

Figure 1. A data recording system flow.

J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*. Copyright © 1999 John Wiley & Sons, Inc.

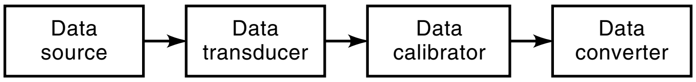

Figure 2. A data acquisition system flow.

be used to sense the temperature that will yield a sequence of electric potential data values. This may be followed by proper calibration and finally the data may need to be converted to a discrete sequence of data using an analog-to-digital converter.

Data Telemetry

When the data source is at a distance from the actual storage medium, then data need to be transmitted through this distance. This technique of distant measurement of data is called data telemetry. The problem here is to get the data from one point to the other without having to go and physically retrieve them. With the progress of satellite communication and internet distribution, this is becoming a wider application area day by day. The Global Positioning System (GPS) is the most active application area of data telemetry. Automated Vehicle Locator (AVL), interactive monitoring and control of remote equipment and processes, stolen vehicle recovery systems, and interactive vehicle messaging systems are some of the applications of GPS.

DATA COMPRESSION

Consider a sequence of NTSC (National Television Standards Committee) standard satellite video data. Each frame when represented in 24-bit true color takes up approximately 1 MB of storage space. With 30 frame/s, 1 min recording of the satellite data thus requires 1800 MB (1.8 GB), which is sufficient enough to fill up the entire hard disk space of a modern day personal computer. Is there a way to reduce the size of these image data? That is what data compression is all about. Based on the requirements of reconstruction, there are two types of compression: lossy and lossless compression. Lossless compression is generally used for text, computer-generated data, and some types of image and video information. A particular application is the compression of bank account and credit information. Lossy compression on the other hand can tolerate some loss of information. For example, a reconstructed video image sequence from a compressed sequence may be well within the tolerance of the viewer.

Lossless Compression

Lossless compression, in general, employs some coding techniques to reduce the size of data.

Huffman Coding. Consider a coding system where the symbols that occur more frequently will have shorter codewords than the symbols which are less likely to occur. The result is a variable-length coding having smaller average code size. In the context of a 256-level gray image, each pixel is coded as an 8-bit code. Now assume that in analyzing the probabilities of the gray levels it is possible to have an average code size of 5 bits per pixel. That gives a compression ratio of 1.6.

Arithmetic Coding. This is used for shorter codeword sources and instances where the codewords have highly skewed probabilities.

Dictionary Coding. A popular realization of this is the LZW approach, which is a variation of the seminal work by Ziv and Lempel in 1977. A dictionary codeword is created from the patterns of the source. The codewords consist of the index, length, and symbol of a frequently used pattern. The LZW technique is used in industry standard applications as in the UNIX `compress` command, GIF (graphics interchange format) image coding, and the V.42 bis modem data compression.

Differential Encoding. Instead of directly encoding a source, the next data element is predicted from its neighboring elements and then the difference between the actual and predicted model is encoded. This reduces the length of code size by reducing the dynamic range of possible values. The basic differential encoding is called Differential Pulse Coded Modulation (DPCM). DPCM attempts to minimize the variance between the predicted model and the actual values. The popular JPEG (joint photographic experts group) technique uses differential encoding for compressing image data.

Lossy Compression

Lossy compression reduces the data size with an acceptable loss of resolution of data. The technique of quantization is in general used for lossy compression, which is then followed by coding. Quantization maps a larger dynamic range of the original source into a smaller range of quantized values. The compression ratio now is a function of the levels of quantization desired, which obviously is much higher than that of lossless compression. The set of quantizer output values is called the codebook.

Differential Encoding. As in the case of lossless compression, this technique may also be used for lossy compression. Here, the differentials of input values go through the process of quantization. Different variants of differential encoding can be found in the literature.

Subband Coding. Here, a data sequence is first decomposed into different bands or resolutions. A bit allocation technique is then used for an optimal combination of these bands. This technique may be used in both lossless and lossy compression. Wavelet decomposition is an example of subband coding. Different types of filters used for subband decomposition result in different types of information.

Transform Coding. The data sequence is first converted into another sequence using some transform. Quantization and other compression techniques are then used in the transform domain instead of in the original data domain. One of the goals of transformation is to wind up with a smaller number of data samples that are representative of the original data sequence. This process is called decorrelating the input sequence, which means that the sample-to-sample correlation of the original data sequence is zero. The transform samples that do not contain much information are discarded. The reverse transformation is then applied on this compressed sequence for reconstruction. There are different types of transforms available, all of which are, to some extent, data dependent. The Karhunen–Loeve transform is an optimal transform in the sense that it minimizes the geometric mean of the variances of the transform coefficients, thus resulting in an optimal decorrelation providing the largest coding gain of any transform coding method. Discrete cosine transform (DCT), discrete sine transform, and discrete Walsh–Hadamard transform are some of the other useful transforms.

Transform coding operates in three steps: First the transformation is applied. The second step involves quantization of the transformed sequence into the desired number of bits. Finally, some types of coding technique like fixed-size code, Huffman, or arithmetic coding are used. The JPEG standard is one of the most widely used techniques for lossy compression of images. It uses DCT as the transform and goes through the following steps.

1. The input image is first level-shifted. This is done by subtracting 2^{M-1} from the image having pixel values in the range $\{0, 2^M - 1\}$. Here, M is the number of bits used to represent each input pixel.
2. Block transformation having a block size of 8×8 using DCT is then applied on the level-shifted image.
3. Quantization on the transform coefficients is then done using variable step size. The idea is to use small step size for low-frequency components and large step size for high-frequency components.
4. Huffman coding method is then applied on the quantized coefficients. Note that a large number of 0 values is obtained after the quantization process and these values do not need to be coded.

As mentioned earlier, the compression ratio really depends on the fidelity of the reconstructed image. In general, a ratio of 25 : 1 results in an acceptable reconstruction for typical image applications.

The MPEG (motion picture experts group) standard, which is used for the compression of motion pictures, also uses the DCT transform. Here the next frame is predicted from the current frame. DCT is taken for the difference of a frame and its predicted frame.

DATA STORAGE

The ever-increasing demand for data storage coupled with the multimedia requirements for text, image, video, and audio storage are growing at an exponential rate. In practice, we encounter three main types of storage: primary or on-line, secondary or near-line, and archival or off-line memory. The range of on-line, near-line, and off-line data storage systems may consist of a variety of technologies such as the magnetic tape drives, magnetic disk drives, optical tape, electron trapping optical memory (ETOM), holographic optical storage, and optical disk drives such as the compact disk read-only memory (CD-ROM), CD recordable (CD-R), digital versatile disk (DVD), and write-one read-many (WORM).

Storage Modes

Data storage may be analog or digital. The audio tape is an example of analog recording, whereas a CD is an example of digital recording. An analog signal is a continuously varying smooth signal having no discontinuity in it. A digital signal, on the other hand, consists of a sequence of discontinuous pulses. The amplitude of the digital signal is fixed, whereas amplitude of analog signal is a changing parameter that imposes a requirement on the dynamic range of an analog storage device. Figure 3(a) shows an analog signal along with discretization in 15 samples. Figure 3(b) shows how each discretized value is represented by using 4 bits. Compared to analog storage, digital storage is more noise tolerant, robust, and easily reproducible. In addition, due to the recent development in the areas of digital technology, control, and communication, digital storage is becoming a more cost-effective solution as well.

Storage Media

Depending on the media and technology used, storage devices may be classified as semiconductor storage, magnetic storage, optical storage, and magneto-optical storage.

Semiconductor Storage. In a computer system, semiconductor storage is used as a faster access medium, which is called primary storage. This is primarily divided into read–write memory (RWM), and read-only memory (ROM). RWM historically is called random access memory (RAM), although ROM is also random access. Random access, in contrast with sequential access of a tape, means any piece of data can be accessed randomly without accessing other data items. RAM is volatile in the sense that its contents is destroyed when the power is turned off. Depending on the technology used, RAM can be static or dynamic. Static RAM (SRAM) does not need any refreshing and is usually faster than dynamic RAM (DRAM). DRAM on the other hand can be made denser and is thus less costly. One variant of ROM called programmable ROM (PROM) can be programmed only once and is sometimes

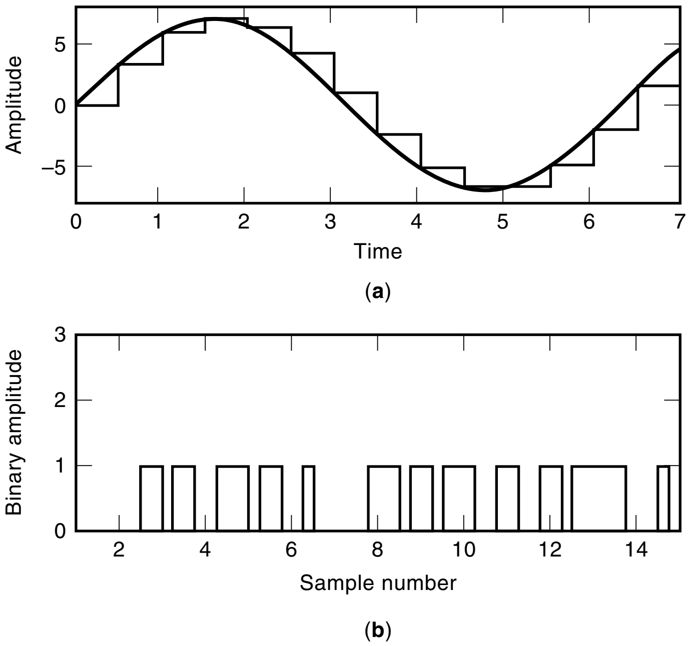

Figure 3. (a) Analog signal and the digitization process. (b) Coding as a digital signal.

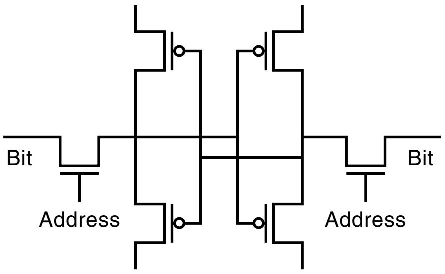

Figure 4. CMOS realization of 1-bit storage of static RAM.

also referred to as one-time programmable (OTP). Erasable PROM (EPROM) can be programmed and reprogrammed many times. EEPROM is electrically erasable PROM that can be erased by exposing it to an electrical charge. This is different from Flash memory in the sense that, while EEPROM stores and erases one byte at a time, Flash memory does it in a block of bytes at a time.

The basic building block of a semiconductor storage system is called a flip–flop. Figure 4 shows the storage of 1 bit of information using a Complementary Metal Oxide Semiconductor (CMOS) SRAM cell. Note that many variations of this configuration are available resulting in different storage capability. With the tremendous success of very large-scale integration (VLSI) technology, semiconductor memory has become the storage device with the highest density and fastest speed. Following are some of off-the-shelf semiconductor storage devices.

PC Card. PC cards can be either storage or input/output (I/O) cards, which is defined by the internationally accepted PCMCIA (Personal Computer Memory Card International Association) standard. PC cards use little power and are compact, highly reliable, and lightweight, making them ideal for battery-powered notebook and palmtop computers, handheld personal digital assistants (PDAs), and personal communicator devices. Their low power, small size, and portability facilitate the use of memory cards in a wide variety of new applications, such as electronic cameras, voice recorders, portable phones, and in the area of data telemetry in general.

ROM and OTP Card. Because the information stored on a ROM card is inseparable from the hardware, a ROM card usually is device-dependent and will not work in a computer system other than the one it is specifically manufactured for. By contrast, OTP cards are built *blank.* After manufacture, these chips can be written to with a special device. While they have the advantage of providing a form of nonvolatile storage, the biggest disadvantage is that once information is stored on the card, it cannot be altered in any way. In addition, the custom manufacturing involved in producing ROM chips makes ROM cards expensive unless produced in large quantities.

SRAM Card. SRAM cards have an advantage over ROM and OTP cards: users can write, modify, or delete the information stored on them. However, SRAM chips are volatile. Affordable SRAM cards are available in capacities of up to 4 MB. Beyond that, the cost of SRAM is prohibitive. Therefore, SRAM cards are used primarily for low volume storage.

Flash Card. One of the latest innovations in memory cards is the use of flash memory technology. Flash cards combine the best features of both ROM/OTP and SRAM cards. Like ROM and OTP chips, flash components are nonvolatile—they retain data even without a power source. And, like SRAM, users can write, modify, or delete the information stored on them. However, use of flash memory currently is prohibiting due to their cost and device capacity.

Magnetic Storage. Magnetic storage exploits electromagnetism. When an electric current flows through a conductor, it generates a magnetic field around that conductor. This magnetic field can influence the magnetic material in the field. When the direction of current reverses, the polarity of the magnetic field is also reversed. So, particles of the magnetic substance can be polarized magnetically in one of the two directions with an electromagnet. Thus, magnetic polarization can be used to distinguish “0” and “1.” This two-way operation of electromagnetism makes it possible to record data on a disk or tape and read the data later. Magnetic disks and tapes are the two variants of magnetic storage devices.

Magnetic Disks. Magnetic disks, which were preceded by magnetic drums, are of two major types: namely, hard disks and floppy disks. Hard disks cannot be bent or flexed—hence the term hard disk. Also, the platters of the hard disk are not removable, and for that reason it is also called a fixed disk.

Construction and Geometry. Disks consist of circular platters made of glass, plastic, or metal and coated with a magnetic substance. A driver motor rotates the disk platter about its central axis. The disk drive also has head motors, which cause head rotation. Heads move radially in and out across the disk surface. Heads are connected to head motors through a moving arm. Usually there is one head per surface. Figure 5 shows the components of a typical magnetic disk drive.

Figure 6 shows the logical geometry of the platters of a disk. Data are stored in concentric tracks on the surfaces of each platter. The number of tracks is an indicator of the storage capacity of the disk. Track density, which is the number of tracks per unit of radial length, is expressed in tracks per millimeters or tracks per inch. A typical disk drive can have more than 2000 tracks per inch (TPI) on its recording surface. A pair of tracks on each side of the disk with the same track number is called a cylinder. In the case of disk packs several tracks may constitute a cylinder. Each track is divided into sectors, which can store a data block. The process of organizing the disk surface into tracks and sectors is called formatting. In almost all systems, including PCs and Macintoshes, sectors typically contain 512 bytes of user data plus addressing information used by the drive electronics.

In earlier hard drive designs, the number of sectors per track was fixed and, because the outer tracks on a platter

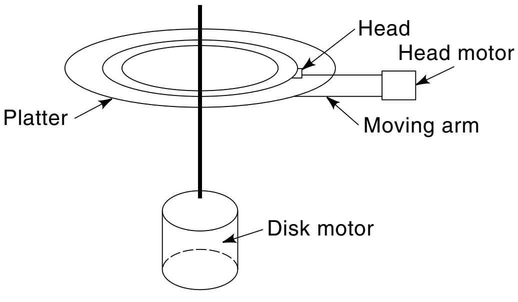

Figure 5. Magnetic disk drive components.

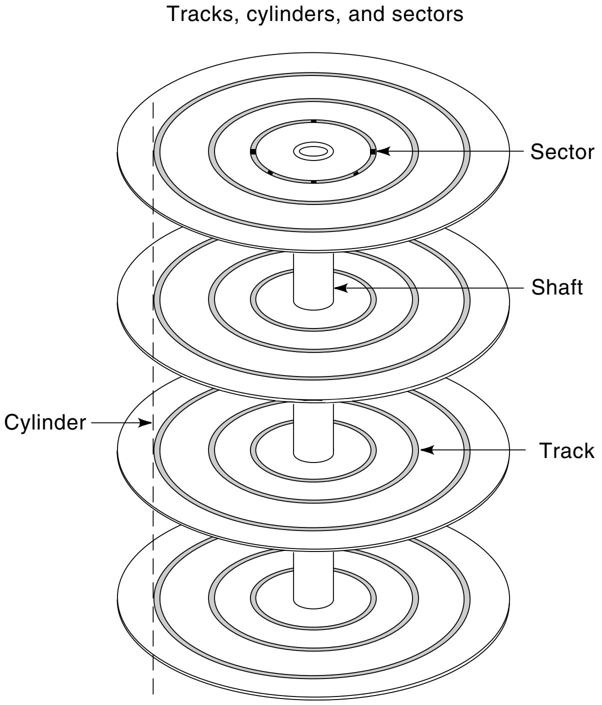

Figure 6. Logical components of a disk system.

have a larger circumference than the inner tracks, space on the outer tracks was wasted. The number of sectors that would fit on the innermost track constrained the number of sectors per track for the entire platter. However, many of today's advanced drives use a formatting technique called multiple zone recording that allows the number of sectors per track to be adjusted so more sectors are stored on the larger, outer tracks. By dividing the outer tracks into more sectors, data can be packed uniformly throughout the surface of a platter, disk surface is used more efficiently, and higher capacities can be achieved with fewer platters.

Read/write heads are the single most important component of a hard disk drive, and their characteristics have a great impact on drive design and performance. A head is a piece of magnetic material, formed almost in the shape of a "C" with a small opening or gap. A coil of wire is wound around this core to construct an electromagnet. In writing to the disk, current flowing through the coil creates a magnetic field across the gap that magnetizes the disk-coating layer under the head. In reading from the disk, the read/write head senses an electronic current pulse through the coil when the gap passes over a flux reversal on the disk.

Several different types of read/write heads exist. Among the earliest are monolithic ferrite heads or those made of a single block of ferrite, a magnetic ceramic material. An improvement on all-ferrite heads were composite heads consisting primarily of nonmagnetic material with a small ferrite structure added. Next came metal-in-gap, or MIG, heads with very thin metal layers added inside the gap to improve magnetic performance. Currently, many drives use thin-film heads, whose name reflects the fact that their structural elements are deposited on a substrate in much the same way that microchips are manufactured. Thin-film technology allows head vendors to achieve much smaller physical dimensions and to better control the fabrication process, both of which result in higher performance products.

Storing and Retrieving Data. The mechanism of storing and retrieving data is almost the same for both the hard disk and floppy disk. The only difference is that a floppy disk has only one disk platter and for this, the number of surfaces is two; the top and bottom surface of the single platter. Most hard disks contain several disk platters, all mounted on the same axis.

When data are retrieved from a hard disk drive, a command is issued to open an existing file, and the application program that is running prompts the user to enter the name of the file to open. It then passes the file name to the operating system, which determines where the file is located on the disk drive—the head number, cylinder, and sector identification. The operating system transfers this information to the disk controller, which drives an actuator motor connected to the actuator arm to position the heads over the right track. As the disk rotates, the appropriate head reads the address of each sector on the track. When the desired sector appears under the read/write head, the entire contents of the sector containing the necessary data are read into the computer's main memory.

Storing data on a hard drive is a similar process to retrieving data, only reversed. The host computer operating system is responsible for remembering the addresses for each file on the disk and which sectors are available for new data. The operating system instructs the controller where to begin writing information to the disk. The controller moves the read/write heads to the appropriate track and writing begins. When the first track is full, the heads write to the same track on successive platter surfaces. If still more track capacity is required to store all the data, the head moves to the next available track with sufficient contiguous space and writes the data there.

Encoding. Traditionally, drives have used analog peak detection read channels. During a write operation, an analog peak detection read channel converts binary, digital data into an analog signal, which the drive's read/write head uses to cause magnetic flux changes on the platter surface [Fig. 7(a)]. During a read operation, the read/write head detects the magnetic flux changes from the data and generates an analog readback signal, in the form of a wave, that it transmits to the read channel. The read channel analyzes the incoming signal to determine the high/positive peaks and low/negative peaks. Finally, it decodes each of these peaks. The limitation of analog peak detection read channels is that, as data densities increase, the analog signal peaks start to overlap, which causes data degradation. To counter this effect, hard drives employ a data encoding scheme during write operations that separate the analog signal peaks. One of these techniques is called partial response maximum likelihood-(PRML).

By using sophisticated digital coding and filtering techniques, PRML read channels sample the analog signal wave at a number of points, as opposed to just at the high and low peaks, as shown in Fig. 7(b). By taking these samples, PRML read channels can determine the *shape* of the readback signal, and thus can interpret the high and low peaks that represent data bits very accurately. Although using the same read/write heads and media, the use of PRML technology provides a much more space-efficient error detection scheme, which yields a 25% improvement in data bit density while achieving the same, low bit error rate as analog peak detection.

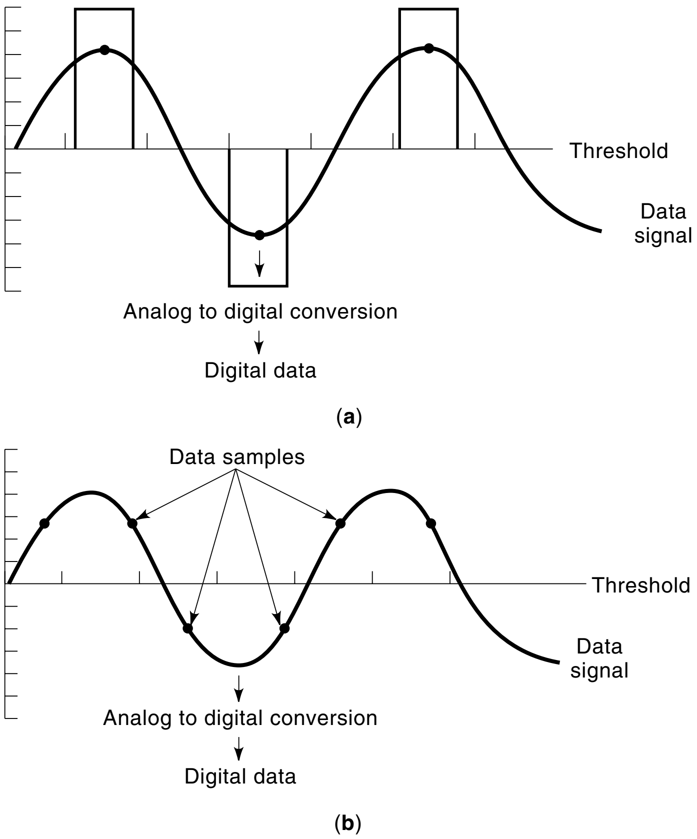

Figure 7. (a) Analog peak detection encoding. (b) PRML encoding.

Recording Codes. Numerous codes for recording are available. Table 1 shows a comparison of them. Code rate is the ratio of number of information symbols (m) and the number of code symbols (n). Code rate $m/n \leq 1$.

Detection or sample window relates to the interval during which the recorded bit may be sampled on playback to determine its value. It is normally measured in terms of bit-cell duration, T. Capacity $C = \log x$, where x is the largest root of the equation,

$$F(x) = x^{k+1} - x^{k} - d - x + 1 = 0 \tag{1}$$

where d is minimum number of zeros between adjacent ones, and k is maximum number of zeros between adjacent ones.

Density ratio (DR) is a measure of the packing density of the code: DR $= m/n(d + 1)$.

Disk Performance Specifications

Data Access Time. Access time is a measure of how long it takes to position a read/write head over a particular track and find the sectors of interest within the track for reading or writing. Thus, data access time is a combination of seek time, head switch time, and rotational latency and is measured in milliseconds.

Seek Time. The amount of time it takes the actuator arm to move the read/write head between tracks is called seek time. For a platter having N tracks on each side, moving from the current track to the next requested bit of data might entail moving just one track or up to $N - 1$ tracks. Seek time is measured in milliseconds (ms). Seek times between adjacent tracks can be as short as 2 ms, while a full-stroke seek (movement between the outer and inner track) consumes about 20 ms. *Average seek time* is defined as the time it takes to position the drive's read/write heads for a randomly located request and usually ranges from 8 ms to 14 ms.

Head Switch Time. The actuator arm moves all of the read/write heads over the platters simultaneously. However, only one of the read/write heads can be reading or writing data at a time. Head switch time measures the average time the drive takes to switch between two of the read/write heads when reading or writing data.

Rotational Latency. Once the read/write head is positioned over the proper track, it must wait for the drive to spin the platter to the correct sector. This waiting period, called rotational latency, is measured in milliseconds and is dependent on the drive's revolutions per minute. On an average, the disk needs to spin half way around before the next sector to be read or written to. Increasing the disk's rotational speed decreases this latency. However, the faster the disk rotation, the hotter the drive runs and the greater the wear on the drive's moving parts. Still, manufacturers are overcoming these problems, and the once universal 3600 rpm (with an average latency of 8.3 ms) is being supplanted by faster speeds like 7200 rpm (with a latency of 4.2 ms).

Data Transfer Rate. Data transfer rate heavily depends on two measures: the disk transfer rate, or how fast data are passed from the disk to the hard drive's controller; and the host transfer rate, or how fast the controller passes data to

Table 1. Comparison of Codes in Magnetic Recording

Code	Rate	Capacity	*d,k*	Packing Density	Detection Window	Dc[a] Free
NRZ-L	1	1	0	1	T	No
NRZ-M	1	1	0	1	T	No
E-NRZ	7/8	1	0, 13	0.875	T	No
FM	1/2	0.6942	0, 1	0.5	T/2	Yes
MFM	1/2	0.5515	1, 3	1	T/2	No
Mil'r Sq.	1/2	0.6509	1, 5	1	T/2	Yes
ZM	1/2	0.5515	1, 3	1	T/2	Yes
RLL 2,7	1/2	0.5174	2, 7	1.5	T/2	No
GCR 4/5	4/5	0.8792	0, 2	0.8	0.8T	No

[a] Dc Free means average value is zero.

the computer's central processing unit (CPU). The data transfer rate is measured in megabytes per second (MB/s). To speed the host transfer rate and minimize mechanical delays (from seeking and rotational latency), manufacturers have added cache buffers to the hard drive. A cache buffer built into the drive electronics provides a temporary storage space that eliminates or reduces any bottleneck in the flow of data between the disk and CPU. A cache buffer coupled with an effective caching algorithm can improve the effective data transfer between a drive and the CPU by a significant factor—often 50% to 100%.

Data Throughput Rate. Data throughput rate is a measure reflecting both data access time and data transfer rate. It represents the total amount of data that the CPU can access in a unit of time. As such, it is a fairly comprehensive measure that factors in most of the major drive performance measurements.

Storage Capacity. Capacities are stated in two different ways: unformatted and formatted. Formatted capacity represents the real capacity, or the amount of data that a user can store; whereas unformatted capacity represents a higher number usually used in advertisements. The unformatted capacity is the maximum number of bytes that can be recorded assuming that each track is continuously and completely recorded with zeros and ones. The unformatted capacity of a disk is, therefore, the product of the maximum number of bytes per track times the number of tracks times the number of recordable sides. For modern recording modes like modified frequency modulation (MFM), one bit is represented by one flux transition. The number of flux transitions for a track equivalent to one complete revolution, that is 2π radian, is therefore calculated as follows:

$$\text{Flux transitions/track} = \text{flux transitions/rad} \times 2\pi = \text{bits/track}$$

The unformatted capacity, C_u, of the 2 Mbytes 3.5 in. diskette (15,916 flux transitions/rad, 80 tracks, 2 sides) is: $C_u = (2\pi \times 15{,}916 \times 80 \times 2)/8 = 2{,}000{,}063$ bytes ≈ 2 Mbytes.

For a certain standardized format, the real formatted capacity is the product of the number of data bytes per sector times the number of sectors per track times the number of tracks per side times the number of sides. For the same diskette using the standardized MS-DOS format, the formatted capacity, C_r, is

$$C_r = 512 \times 18 \times 80 \times 2 = 1,474,560\,\text{bytes} \approx 1.47\,\text{Mbytes}$$

For this type of diskette the formatted capacity is about 73.7% of the unformatted capacity.

The performance of magnetic storage devices has been enhanced manifold by emerging technologies. Two of these are mentioned below. One relates to capacity and the other to speed.

- Materials such as thin-film and magnetoresistive (MR) heads have led to improvements. Unlike current head technologies—all of which are basically tiny inductive electromagnets—MR technology uses a different material whose electrical resistance changes in the presence of a magnetic field. A small stripe of MR material is deposited on the head structure and, as it passes over the magnetic patterns on the disk, it senses the strength of the magnetic field and creates electrical pulses corresponding to the flux reversals. This mechanism cannot be used for writing, so a conventional thin-film inductive write element is deposited alongside the MR stripe. MR head technology began to appear in drive designs in 1994, with wider acceptance and inclusion in more designs by 1995. IBM recently announced its MR head-based disk that holds 3 billion bits of data per square inch.
- Seek command reordering of the disk drive is another improvement. In the absence of command reordering techniques, the disk drive would execute the commands in the queue in the order it received them. The result is that the read/write head randomly sweeps over the surface of the platter executing each command in the queue in sequence. Command reordering minimizes this random movement of the read/write head. One technique for command reordering is called *shortest seek time-first* ordering that selects the command with shortest seek time for next execution. The problem with the method is that it gives higher priority to the tracks located near the middle compared to the commands requiring data from the outer and inner tracks because their seek times are usually longer. *Elevator seeking* overcomes this problem by also ordering commands such that the read/write heads move back and forth across the entire platter surface. Optimized reordering command algorithm (ORCA) takes command reordering one step further to calculate the optimum sequence of seeks so as to minimize the combined total of seek time as well as rotational latency. Thus, it examines both the track location of the seek on the disk and the sector location of the seek on the track. It orders the seek command tasks in a more efficient sequence, increasing the data throughput of the drive by an average of 20%.

Magnetic Tape. Commonly used on early mainframe computers, one of the first computer storage technologies was the magnetic tape drive, or simply tape drive. Magnetic tape is sequential data storage device. To read data, a tape drive winds through the spool of tape to the exact location of the desired information. To write data, the tape drive encodes it sequentially on the tape. Because tape drives cannot randomly access or write data like disk drives, and are thus much slower, they have been replaced as the primary storage device in most computer applications. However, with its high storage capabilities and low cost-to-megabyte ratio, magnetic tape is still very much in use as a storage medium for archiving large amounts of data. Recent advances in tape technology, such as faster tape drives and digital audio tape (DAT) cartridges, have also made tape a preferred technology for backing up network servers and other critical data.

Optical Storage. Currently, optical memory devices are widely used as secondary memory in computer systems. The most widely known optical storage device is CD-ROM, which has significantly enhanced the replication, distribution, and utilization of software, games, audio, and video intended for use in computers or in entertainment instruments. For instance, a CD-ROM containing 680 Mbytes of data can be mass duplicated by injection molding in a few seconds costing less than 10 cents per piece. The main attractive features of an optical data storage system are the removability, portability, and reliability.

Optical Disk. An optical disk is usually constructed from a plastic or glass substrate coated by one or more thin-film layers and contains a number of tracks along which information is stored. The manufacturer may record the information on the substrate surface or information may be recorded by the user on one or more of the thin-film layers along the tracks. The storage layer is generally sandwiched between two dielectric layers and the stack is covered with a metal layer for better reflectivity. The functions of the dielectric and reflective layers include reduction of thermal cross talk during writing, optimize absorptivity and reflectivity, and protecting the storage layer.

Figure 8 shows the basic building blocks of an optical disk system. The light beam emanating from the laser is collimated and directed toward a high-numerical-aperture objective lens via a beam splitter. The objective lens focuses the light beam on the surface of the disk where information is written to or read from a given track. The light reflected from the disk surface is collected by the objective lens and directed toward one or more detectors through the beam splitter. The detector extracts the data readout signal as well as the focusing and error-tracking signals.

The main features of an optical disk are shown in the right side of Fig. 8. The access window is used by the read/write head to gain access to the disk. The hub provides the mechanical interface with the drive for mounting and centering the disk on the spindle. The track shown in Fig. 8 is made up of concentric rings of width W_t and adjacent tracks are separated from each other by a guard band of width W_b. If the binary digits (0 and 1) are stored as transparent and opaque marks of length L_m and the entire disk surface of radius, r, is covered with such marks, then the capacity of the optical disk may be expressed as

$$C = \frac{\pi r^2}{(W_t + W_b)L_m} \tag{2}$$

where C is expressed in bits per surface. For a typical 5.25 in. optical disk, r = 67 mm, $L_m = 0.5$ μm (determined by the read/write laser wavelength), and $W_t + W_b = 1$ μm. The capacity of this disk, using Eq. (2), computes to approximately equal to 3.5 gigabytes per surface, which is pretty close to reality.

In optical disk and tape storage systems, semiconductor laser diodes having the shortest possible wavelength that can provide enough power for read, write, and erase operations for several thousand hours are generally used. In general, the size of the focused spot varies directly with the wavelength. Therefore, shorter wavelength lasers can be focused to smaller spots at the diffraction limit. The diameter d of the focused spot may be expressed as

$$d \approx \frac{\lambda}{NA} = \frac{\lambda}{\sin\theta} \tag{3}$$

where NA refers to the numerical aperture of the objective lens, λ is the wavelength of light, and θ is the half angle subtended by the focused cone of light at its peak. Equation (3) implies that higher numerical aperture objective lens should be used for smaller spots. But the depth of focus δ of the objective lens may be expressed as

$$\delta \approx \frac{\lambda}{(NA)^2} \tag{4}$$

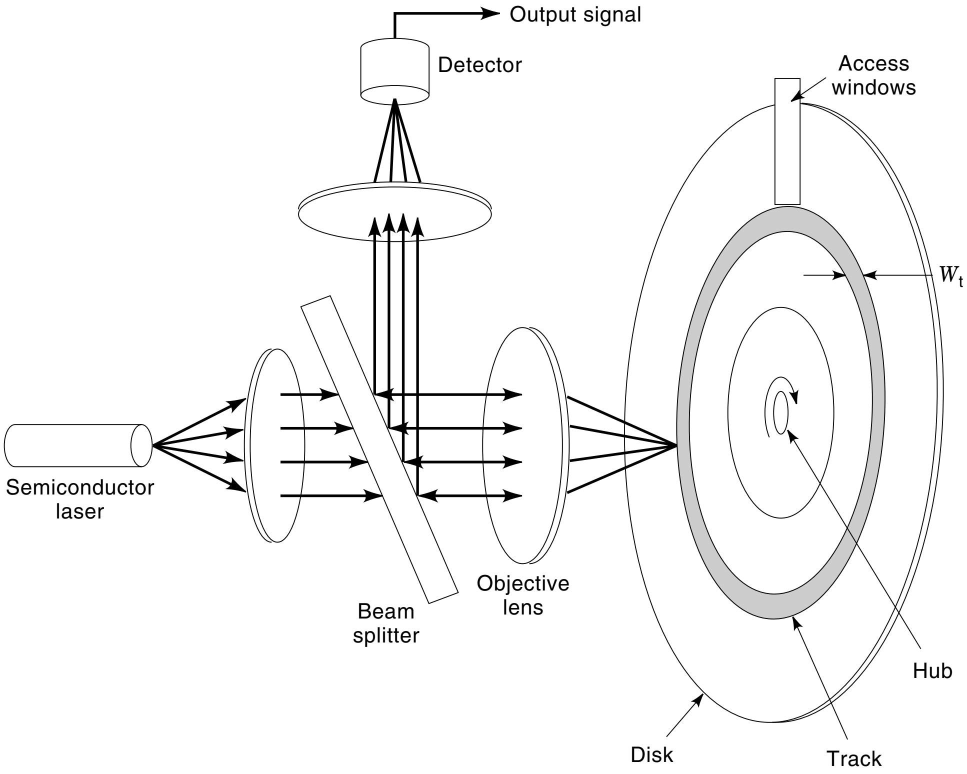

Figure 8. Basic building blocks of an optical disk system: W_t is width of track.

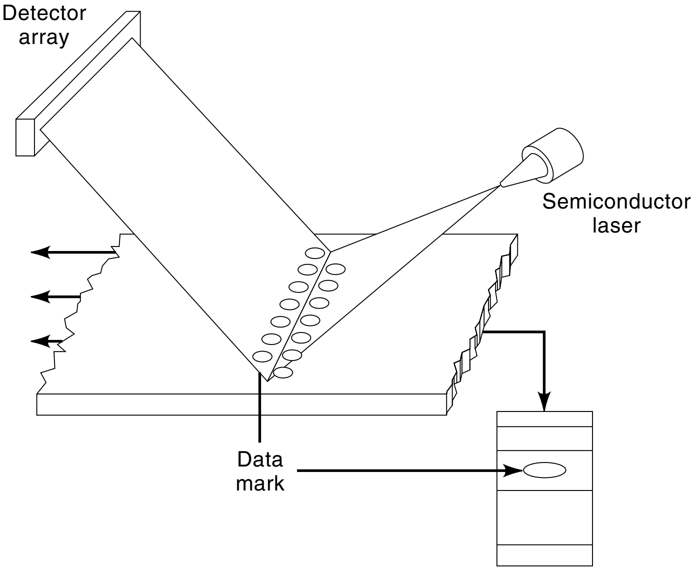

Figure 9. Basic building blocks of an optical tape system.

insinuating that the depth of focus will be smaller for higher numerical aperture objective lens. The range of *NA* usually varies from 0.4 to 0.6.

Optical Tape. Figure 9 shows the main components of an optical tape which consists of a substrate onto which an active layer and a dielectric film are sputtered flanked by a back and a top coating for better reliability. The active layer usually consists of thin metallic film or dye polymer or phase change material. Information is hole punched (called marks) into the active layer by a focused semiconductor laser beam. For data read out, the tape is scanned by a focused laser beam and the marks cause the amplitude of the reflected light to vary. These variations are detected by a detector array to generate the appropriate data readout signal.

Optical tape is mainly used for high volume storage. It is so thin that very large amounts of data can be stored in a very small volume of it. For instance, optical tape has at least 25 times the volumetric density of CDs. Although magnetic tape is economical, it is not suitable for long-term archiving because the magnetic medium is degraded by creep, track deformation, and print-through, requiring the transfer of stored information to a new magnetic tape after a given interval. Optical tape is more durable, does not suffer from print-through, and is more reliable than the magnetic tape. The effective data rate for a magnetic tape is very low since the read/write process is inherently sequential. On the other hand, the effective data rate for an optical tape can be enhanced significantly by exploiting parallel read/write channels as depicted in Fig. 9.

Electron Trapping. Electron trapping optical memory (ETOM) exploits multilevel encoding through amplitude modulation which allows it to store one of several intensity levels at a given location. For instance, if any of eight amplitude levels can be detected reliably, then 3 bits can be stored in a single location, thus tripling the storage capacity.

ETOM usually consists of an active layer of phosphor-like material coated on a substrate. To store information, the trapping material is exposed to light of one wavelength, which excites electrons into the conduction band where they are trapped. The number of electrons trapped is directly proportional to the intensity of light. To read data, the trapping material is exposed to light of another wavelength, which causes the conduction band electrons to drop by emitting photons of a third wavelength. The number of photons emitted is proportional to the number of electrons trapped initially. Thus, by counting the number of photons multiple values can be stored at a location.

The main limitations of an ETOM are that it requires a complex read/write head and three different lasers. Furthermore, adding an extra bit for a given location requires twice as many decibel levels, thus causing the loss in signal-to-noise ratio.

Holographic Storage. Figure 10 shows the general features of a holographic data storage system. A hologram is created by recording the spatial interference pattern of two coherent beams, called the object beam and the reference beam, which are usually generated from the same laser source using a beam splitter and mirror combination. The object beam carries the information to be recorded while the reference beam is just a plane wave. The information or the data to be stored is usually a two-dimensional arrangement (called pages) of transparent and dark patterns representing binary ones and zeros. This two-dimensional pattern is displayed in a spatial light modulator (SLM) in the data plane where it is modulated onto the object beam. The object beam interacts with the reference beam in the holographic material to form a fringe pattern, which is recorded by the material, thus forming the hologram. When the hologram is illuminated with the original reference beam, it reconstructs the object beam, thus retrieving the information stored in the hologram.

As the hologram recording material becomes thicker, the reconstruction becomes particularly sensitive to the angle between the two interfering beams. Therefore, multiple holograms may be superimposed in the same volume by changing the angle of incidence of one of the beams: that is, the reference beam. This is accomplished by employing another SLM called the reference plane SLM. Each hologram can be accessed independently by using an appropriate set of reference beams.

The Bragg effect governs the diffraction from a hologram, which implies that if the reference beam is applied to the holographic medium at an angle other than that used during recording, it will not diffract the stored fringe pattern or the hologram. Therefore, thousands of holograms can be recorded in a small volume hologram (on the order of 10^{11} bit/cm^3) using a reference beam with a Bragg angle unique to each of the stored patterns. Furthermore, the page-oriented parallel access to data in a hologram offers low access times (less than 1 ms) and potentially high data rates (gigabit/s).

Magneto-Optical Storage. Magneto-optical (MO) disk systems combine the technology of traditional magnetic media, like hard disk drives, with optical disk technology. MO technology allows users to pack hundreds of megabytes of data on a disk that looks similar to a traditional 3.5 in. floppy disk and typically comes in a 3.5 in. or 5.25 in. form factor. An MO disk is made of materials that cause it to be highly resistant to magnetic fields, or coercive force, at room temperature.

An MO drive writes to the disk using a read/write head assisted by a laser. A laser heats up the disk surface to its Curie point. Then, the read/write head passes over the disk,

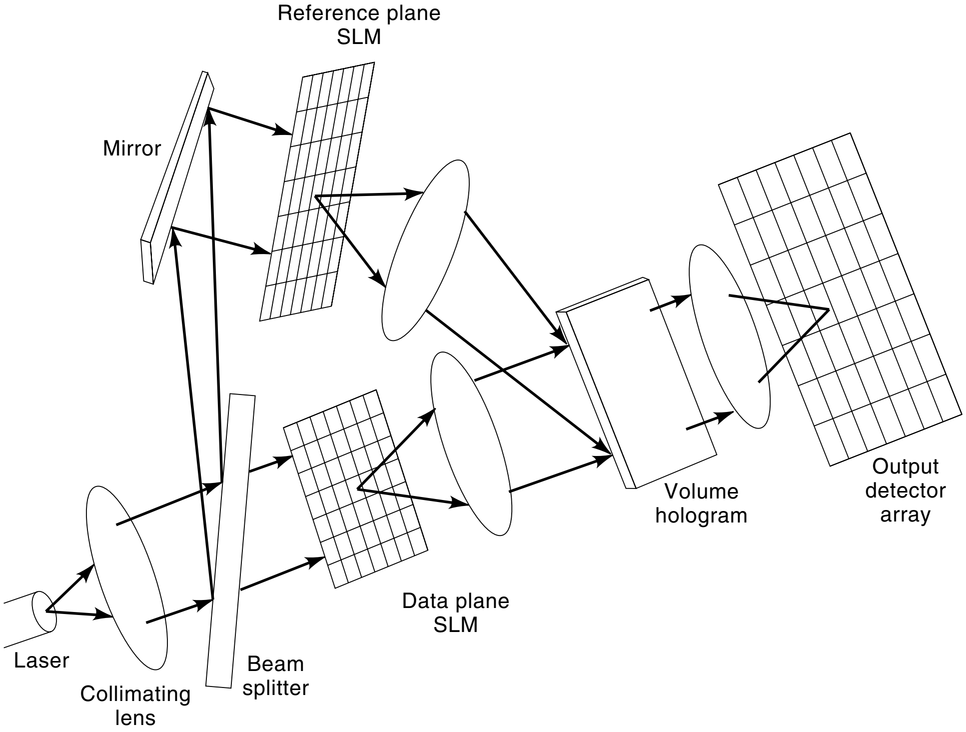

Figure 10. Basic building blocks of a volume holography storage system.

polarizing those areas that have been heated by the laser. Because a laser can be focused on a much narrower field than a traditional, magnetic read/write head, data written on an MO disk with the assistance of a laser results in a very high data density not available with traditional hard disk drive technology. During a read operation, the MO drive uses the same laser to sense the data stored on the disk. As the laser scans the disk surface, the drive detects a reflection of light by the data bits oriented in one direction and no reflection from the data bits oriented in the opposite direction. Thus, an MO drive can distinguish between the 0 and 1 data bits stored on the disk. MO disks have many advantages:

- They provide very high data densities, achievable because of the use of a laser.
- Data can be changed at will—added, modified, or deleted.
- Data are resistant to magnetic fields. Unlike a traditional floppy or hard disk, a magnetic field alone cannot alter the data without the additional heat provided by a laser.
- Because of the use of the laser to assist in reading and writing data, the read/write head does not need to be as close to the surface of the disk as with a hard disk drive.

The disadvantage of MO technology is that, because of the high intensity of the magnetic field created with the combined use of the read/write head and laser, the drive cannot change the field polarity very quickly. Therefore, the drive must pass over the disk twice to write to it. Of course, attempts are going to compress the writing process into a single disk rotation. So, MO drives with one-pass writing capabilities should be available soon.

There are many other devices used for optical storage. Multiple quantum well (MQW)-based self electro-optic effect devices (SEED) operate by changing their optical absorption characteristics inside the quantum-sized media. Nonlinear devices such as photorefractive materials use techniques like phase-conjugation and two-beam coupling for optical storage.

BIBLIOGRAPHY

1. C. D. Mee and E. D. Daniel, *Magnetic Recording,* New York: McGraw-Hill, 1988, Vols. I–III.
2. J. Watkinson, *An Introduction to Digital Audio,* 1st ed., Boston: Focal Press, 1994.
3. G. Held, *Data Compression: Techniques and Applications Hardware and Software Configurations,* 1st ed., New York: Wiley, 1984.
4. J. A. Storer, *Data Compression: Methods and Theory,* 1st ed., Rockville, MD: Computer Science, 1988.
5. K. Sayood, *Introduction to Data Compression,* 1st ed., San Mateo, CA: Morgan Kaufman, 1996.
6. M. H. Kryder, The special section on magnetic information storage technology, *Proc. IEEE,* **74**: 1475–1476, 1968.
7. F. Jorgensen, *The Complete Handbook of Magnetic Recording,* Blue Ridge Summit, PA: Tab Books, 1980.
8. C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech J.,* **27**: 379–423 (part 1), 623–656 (part 2), 1948.
9. N. D. Mckintosh and F. Jorgensen, An analysis of multilevel encoding, *IEEE Trans. Magn.,* **MAG-17**: 3329–3331, 1981.
10. G. V. Jacoby, Ternary 3 PM magnetic recording code and system, *IEEE Trans. Magn.,* **MAG-17**: 3326–3328, 1981.
11. M. Camras, *Magnetic Tape Recording,* New York: Van Nostrand Reinhold, 1985.
12. K. Ford (ed.), *High Density Digital Recording,* NASA, Scientific and Technical Information Branch, NASA reference publication, 1111, 1985.
13. A. S. Hoagland, *Digital Magnetic Recording,* Malabar, FL: Krieger, 1983.

14. H. Jouve, *Magnetic Bubbles,* Orlando, FL: Academic Press, 1986.
15. Y. Manolopoulos and A. F. Vakali, Seek distances in disk with two independent heads per surface, *Inf. Process. Lett.,* **37**: 37–42, 1991.
16. M. Mansuripur and G. Sincerbox, Principles and techniques of optical data storage, *Proc. IEEE,* **85**: 1780–1796, 1997.
17. G. Bouwhuis et al., *Principles of Optical Disk Systems,* Bristol, UK: Adam Hilger, 1985.
18. A. B. Marchant, *Optical Recording,* Reading, MA: Addison-Wesley, 1990.
19. T. W. McDaniel and R. H. Victoria (eds.), *Handbook of Magneto-Optical Data Recording,* Westwood, NJ: Noyes, 1997.
20. M. Guizani, ATM architectures using optical technology: An overview of switching, buffering, and multiplexing, *Int. J. Network Manage.,* **7** (4): 198–220, 1997.
21. M. A. Karim and A. A. S. Awwal, *Optical Computing—An Introduction,* New York: Wiley, 1992.
22. M. Mansuripur, *The Physical Principles of Magneto-Optical Recording,* Cambridge, UK: Cambridge Univ. Press, June 1995.
23. B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics,* New York: Wiley, 1991.
24. M. M. Murshed and M. Kaykobad, Seek distances in two-headed disk systems, *Infor. Proc. Lett.,* **57**: 205–209, 1996.

FARID AHMED
Penn State Erie, The Behrend College

MOHAMMAD S. ALAM
Purdue University

M. KAYKOBAD
Bangladesh University of Engineering and Technology