

















































































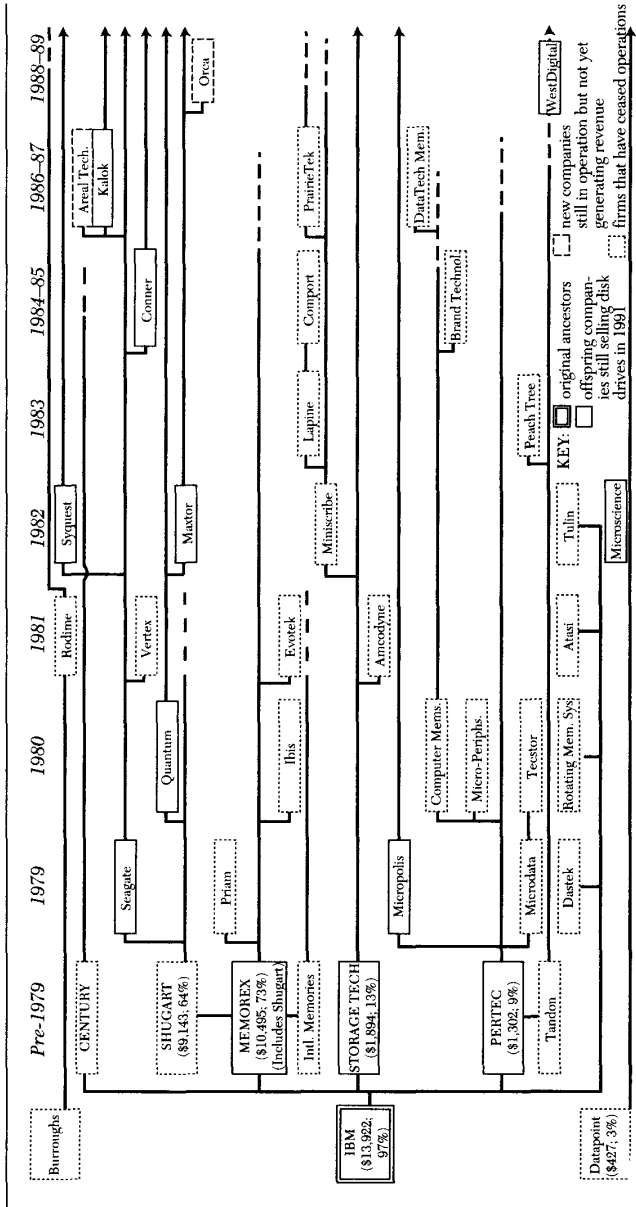








Figure 7  
Employment Genealogies of Founders of Leading Start-Up Disk Drive Manufacturers



dants, in \$ millions, and the percentage of the total for which they accounted are listed in the IBM, Datapoint, Memorex, STC, Pertec, and Shugart boxes.) Moreover, six of the seven Shugart spin-offs that generated revenue were still in operation in 1991 and included the U.S. industry's four largest firms—Seagate, Conner, Quantum, and Maxtor.<sup>38</sup> This compares with one of eight Pertec spin-offs still in operation (Micropolis) and none of the descendants of Storage Technology.

One might term the pattern of spin-outs shown in Figure 7 one of *horizontal dis-integration*, where independent firms split off from predecessor firms to focus on different market segments. There seems to have been a pattern of *vertical dis-integration* in the industry as well. Whereas the industry's original dominant firms were thoroughly integrated into component manufacturing and the research required to support advanced component development, those firms' focus on a single set of (primarily internal) customers induced the spin-out of a host of independent companies that supplied components to the progressively less integrated set of disk drive manufacturers arrayed in Figure 7. Although the story of the creation of a network of independent component manufacturers is complicated enough to merit its own history, a brief summary of how and why independent component manufacturing firms spun out of the industry's initially dominant integrated firms will be helpful here.<sup>39</sup>

Firms that entered the industry with a new architectural technology targeted at an emerging market generally shifted their technological sights toward improved component technology, because improved componentry was the engine of performance improvement within each established product architecture. This shift in technology strategy from architectural orientation to component orientation entailed a significant change in the economics of product development for IBM and the other early industry leaders, because no network of component supply firms existed. Development of new product architectures generally had been an engineering task, not an issue of research and development. But the development of new component technology required substantial investment in research

<sup>38</sup> The \$9.143 billion in cumulative revenues shown in Figure 7 for Shugart's progenitors does not include the approximately \$1 billion in revenues booked by Seagate in 1989 from its acquisition of Control Data's disk drive operations.

<sup>39</sup> An initial version of this history can be found in Christensen, "Industry Maturity and the Vanishing Rationale for Industrial Research and Development."

and development, because component technology development is where basic scientific research, engineering, system design, product design, and process development all come together. In terms of the time, expense, and expertise required, component technology development was an enterprise fundamentally different from the design of new product architectures that employed available componentry. This difference between the pace and scale of component development and those of product design eventually made it impossible for the integrated firms to perform internally the coordinating roles that Chandler observed in other industries in funneling new component technologies into new product designs.

The following account of IBM's development of the thin-film head illustrates four phases that are typical of the course of events that occurred in the development and diffusion of many components within most of the industry's early leading integrated firms, including Control Data, Xerox (which supported its Diablo, Shugart, and Century disk drive divisions through its Palo Alto Research Center), Burroughs, Digital Equipment, and Seagate Technology.

*The Disparate Cycles of Component and Systems Development: IBM and the Thin-Film Head* • As the industry's pioneer and dominant firm during this period, IBM led the shift in technology strategy from architecture development to component development. Improved component technology defended IBM's large, growing, and very profitable mainframe business by providing its customers with steady performance improvement within the 14-inch Winchester architecture. Backward integration reduced the uncertainty that required components might not be available to meet customer and competitive requirements in the next product generation.

In the earliest stages of thin-film head development, IBM worked alone to expand its understanding of basic scientific issues such as the physics of magnetic recording and the properties of new materials. This phase began in 1965 in IBM's advanced research facilities at Yorktown Heights, New York, and San Jose, California. The second phase (beginning roughly in 1971 and ending in 1976–78) was stimulated by proof of the concept at IBM and by the spread of that information through published scientific papers and the trade press to other firms. Statements by respected IBM scientists that thin-film technology was important and feasible led a broader group of vertically integrated manufacturers—Burroughs, Control Data, Digital Equipment, Fujitsu, Hewlett Packard, Hitachi, and NEC—to

initiate their own development efforts. IBM's pathbreaking research resolved a lot of uncertainty for these other firms: once they knew that something could be done, they could focus with greater commitment on learning how to do it. The end of this second phase was marked by the building of early working prototype thin-film heads.

In the third phase, the component design was refined, a manufacturing process was established, and the component was designed into a new disk drive model. In 1976 Burroughs was the first in the industry to announce a drive with a thin-film head, but it was never able to manufacture the head reliably and withdrew the drive from the market. IBM introduced its model 3370, equipped with thin-film heads, in 1979. Positioned at the highest-performance end of IBM's line, the 3370 was a very successful product, even though the heads were extremely difficult and expensive to manufacture.

To this point, the component technology leadership of the large, vertically integrated firms—particularly IBM—was unambiguous. In moving to the fourth phase, however, the story becomes troubled. Although IBM usually initiated each component development process in behalf of the industry and was often the first to introduce the new component in a high-end product model, IBM and the other vertically integrated manufacturers subsequently were very slow to employ the componentry they had developed in other new models in their product lines. The vertically integrated firms' commercial introduction of the new components in a limited number of high-end models typically stimulated the fourth phase in the emergence of new component technology, in which demand for the new componentry became intense among certain independent, non-integrated disk drive manufacturers. These independent firms, such as Maxtor and Micropolis, generally pursued technology strategies that pushed, through innovative (some would say risky) system design, what was called in the industry "the bleeding edge" of performance—a much more aggressive engineering posture than the vertically integrated manufacturers typically were inclined to adopt.

IBM viewed its proprietary access to advanced componentry as its primary competitive advantage and was reluctant to sell its components in the external marketplace, because it viewed the bleeding-edge manufacturers that most needed thin-film heads as indirect competitors.<sup>40</sup> Sensing the opportunity to match component supply

<sup>40</sup> Although none of the IBM engineers or executives interviewed for this history cited antitrust pressure as a force that kept IBM from selling its components in the OE market, Downloaded from <https://www.cambridge.org/core>. IP address: 207.241.231.81, on 03 May 2019 at 03:42:24, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.2307/3116804>

and demand more closely, several venture capitalists recruited key IBM and Xerox engineers into new start-up firms to produce and sell the new-technology components to bleeding-edge disk drive makers in the original equipment market. This was possible because the vertically integrated firms typically enjoyed little patent protection for the components: much of the key technology consisted of process know-how. The industry's leading thin-film disk manufacturer, Komag, and the leading thin-film head manufacturer, Read-Rite, both started in this manner. In 1992, within eight years of their founding, these firms together logged over \$1 billion in revenues.

When the components became available from these start-up firms, even on an irregular, unpredictable, low-yield basis, disk drive manufacturers tended to utilize the new technology throughout their product lines at a pace that was roughly commensurate with their distance from the bleeding edge. And in general, the integrated manufacturers such as IBM tended to design conservatively, some distance from the bleeding edge. As a result, although the vertically integrated firms were the first to develop and introduce the new components, they were the slowest to incorporate them across the breadth of their product lines (see Figs. 8a and 8b). This was most strikingly the case with IBM. Although it spent over \$300 million developing the thin-film head, it was the last firm in the industry to use the technology broadly in its entire product line. And although IBM spent over \$100 million developing thin-film disks, it did not use them in any commercial product until 1988.


The unfortunate combination of the integrated firms' leadership in component development and their followership in component use seems to have been the result of an inexorable decoupling of component development from systems design as the industry matured. This was a problem that IBM worked hard to remedy but that in many ways derives from the fundamentally different natures and time scales of the two processes. Initially, because it was such a dominant presence in the computer markets, IBM was able to control the pace at which it introduced new models into the market—roughly every four years in the 1970s. IBM was therefore reasonably capable of coordinating the emergence of new long-lead-time components with the design of the product systems in which they ini-

---

many outside observers believe that IBM was reluctant to sell key components in the open market for fear of fueling the U.S. Department of Justice's antitrust suit against the company, pending at this time.

Downloaded from <https://www.cambridge.org/core>. IP address: 207.241.231.81, on 03 May 2019 at 03:42:24, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.2307/3116804>

Figure 8  
 Differences in the Rates at Which Vertically Integrated Computer Makers and the Largest Independent OE Market Manufacturers Adopted Thin-Film Heads and Thin-Film Disks across Their Product Lines




tially would be used. As the computer industry matured, however, competitive market forces increasingly became the drivers of the product development cycle. As IBM lost its ability to control that cycle, coordination between component development and system design became inherently more difficult.

When the component development cycle was decoupled from the product development cycle by the competitive market, it simply became impossible for anyone to predict accurately, a decade in advance, whether the company's product position in the market would demand that a specific new component technology be used in its products. The market shifted, the conventional technology progressed, and—possibly the most unpredictable of all—the company's product-market strategy changed. When fully committed development of thin-film head and disk technology was initiated in the early 1970s, almost no one could have imagined that, a decade later, 80 percent of IBM's drives would be used in relatively non-demanding desktop personal computer applications.

In response to the decoupling forces, IBM's managerial practice evolved toward a sort of "free market" system. When new component technologies were developed, they were made available to (but not forced on) product design engineers. The designers of new disk drive systems were free to choose whatever component technologies IBM had in its arsenal to meet the performance objectives of their product. Under this arrangement, IBM's market position—the demands of its customers—determined the pace at which the company employed advanced component technologies in its new models. Figure 9 charts over time the product positions (average megabytes of all models in the 5.25-inch architecture) of Micropolis, Maxtor, and IBM. It shows that in 1984 Maxtor's average 5.25-inch model held 125 megabytes, Micropolis's held 64, and IBM's held 16.5. There is nothing normative about this observation—these firms were simply serving different markets. Maxtor was selling to the memory-starved engineering workstation market, whereas IBM was making 5.25-inch drives for its XT and AT personal computers.

Figure 9 also shows the percentage of each firm's 5.25-inch models in a given year that employed thin-film heads, thin-film disks, and RLL codes. Each of these technologies was developed at IBM but, in 1984, none of them was used in the IBM or Micropolis product lines, whereas Maxtor used thin-film disks in all of its models, with ferrite heads and MFM codes. In 1985, Micropolis adopted thin-film disks on 100 percent of its new models, and Max-

Figure 9  
Capacity Points at which Maxtor, Micropolis, and IBM Incorporated  
Advanced Component Technologies in Their 5.25-Inch Disk Drives



tor and Micropolis converted completely to RLL codes. But IBM still did not use these technologies—it continued to support its market position with established technologies, which were much less costly and risky to use. In 1986, Maxtor and Micropolis both began using thin-film heads, while IBM was still able to satisfy its requirements with conventional technology. Finally, when its average 5.25-inch drive approached the 80–100 MB range in 1987, IBM began using RLL codes in 88 percent of its new models (Micropolis had adopted RLL codes when its products had penetrated this same range two years earlier). When its 5.25-inch drives reached even more demanding territory in 1988, IBM adopted thin-film disks on 62 percent of its new models. Although this step was taken four years after Maxtor had adopted thin-film disks, it occurred when IBM's drives reached the same capacity territory that Maxtor and Micropolis had occupied when they first used thin-film disks. Finally, though IBM had first used thin-film heads in a few high-end 14-inch drives as early as 1979, thin-film heads were not used in its 5.25-inch models until 1990.

By drawing horizontal lines across Figure 9, one can create product performance zones that seem to have mandated the use of particular component technologies. IBM's computer business had devolved to a performance position far from the market's leading edge by the mid-1980s, and it had not yet entered the zones requiring use of the component technologies it had developed. On this basis, it would be difficult to argue that IBM's failure to utilize more broadly the component technologies it had paid so dearly to develop was the result of conservative or inept technical management. Simply put, the firms that needed the new component technologies used them; the firms that did not, did not.

Because IBM turned out not to need all of the component technology it had developed when it became available and yet had a policy of not selling components outside the company, the independent component suppliers that spun out of IBM generally grew to become larger producers of the components than did IBM itself. As a result, the firms that incurred the development costs of the new component technology were not those that generated the revenues (see Table 7 for the case of thin-film heads). IBM, which spent one-third of the industry's R&D dollars for this technology, has produced altogether only 8 percent of the industry's cumulative output of thin-film heads. The independent start-ups, after incurring only 15 percent of development costs, have captured most of the market. A

Downloaded from <https://www.cambridge.org/core>. IP address: 207.241.231.81, on 03 May 2019 at 03:42:24, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.2307/3116804>

*Table 7*  
 Amount and Share of Development Costs Incurred vs. Amount and  
 Share of Total Production of Thin-Film Heads  
 (includes captive use and OE market sales—dollars in millions)

<i>Firm or Group</i>	<i>Total R&amp;D Costs</i>	<i>R&amp;D per Firm</i>	<i>Percent of Industry R&amp;D Costs</i>	<i>Percent of Industry Units Produced, 1990</i>
IBM	\$300	\$300	32	8
Other Vertically Integrated <sup>a</sup>	\$500	\$86	53	38
Independent Start-Ups	\$150	\$15	15	54

<sup>a</sup> Burroughs, Control Data, Digital Equipment, Fujitsu, Hitachi, NEC

Sources: Estimates given by former employees; *Disk/Trend Report*; Peripheral Research Corporation.

similar table could be constructed for thin-film disk technology as well. Although the disparity between costs incurred and units produced may represent an unfortunate subsidization of the competition from the point of view of the IBM shareholders, it is clear that the disk drive industry—and a great many entrepreneurs and venture capitalists within it—have benefited greatly from IBM's extraordinary technological largesse.

### Conclusion

Coordinating the very different enterprises of component development and product system design in an increasingly segmented market became a nearly impossible challenge for the integrated firms as the disk drive industry matured through the 1980s. And in the face of strong, diverse market demand, managers in vertically integrated firms found it difficult to protect or retain valuable component and architectural technologies that their customers did not want. As a result, by the 1990s nearly all of the industry's firms had decoupled their vertically integrated operations to some degree, enabling groups at each stage of the value chain to sell their output in the original equipment market. Control Data, Fujitsu, Hitachi, and NEC had followed a policy of selling completed disk drives in the OE market from the late 1970s. By 1990, IBM, Digital Equipment, and Hewlett Packard all had followed suit, selling disk drives aggressively

Downloaded from <https://www.cambridge.org/core>. IP address: 207.241.231.81, on 03 May 2019 at 03:42:24, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.2307/3116804>

in the OE market to their computer system competitors, rather than continuing to cede that business to defecting engineers. And some firms, such as Digital Equipment and Seagate, became leading suppliers of thin-film disks and heads, not just to their own downstream disk drive operations, but to direct disk drive competitors as well. Hence, an industry whose foundation and growth were built through the activities of large-scale, integrated organizations became in its more mature years an industry where market mechanisms forced the decoupling and specialized focus of enterprises that once were extensively integrated. This process created an industry structure where market mechanisms and interfirm transactions, rather than managerial coordination within large-scale firms, became the means for coordinating the development and manufacture of disk drives.

## Appendix

### How Disk Drives

#### Work and a Glossary of Technical Terms

Disk drives write and read information in the same sort of binary code that computers use. Most disk drives comprise a read-write head, mounted at the end of an arm that swings over the surface of a rotating disk in much the same way that a phonograph needle and arm reach over a record; disks, which are aluminum or glass platters coated with magnetic material; at least two electric motors—a spin motor that drives the rotation of the disk, and an actuator motor that moves the head to the desired position over the disk; and a variety of electronic circuits that control the drive's operation and its interface with the computer. The read-write head is a tiny electromagnet, whose polarity changes whenever the direction of the electrical current running through it changes. Because opposite magnetic poles attract, when the polarity of the head becomes positive, the polarity of the area on the disk beneath the head switches to negative, and vice versa. By rapidly changing the direction of current flowing through the head's electromagnet as the disk spins beneath the head, a sequence of positive- and negative-oriented magnetic domains are created in concentric tracks on the disk's surface. Disk drives can use the positive and negative domains on the disk as a binary numeric system—1's and 0's—to "write" information onto disks. Drives read information from disks in essentially the opposite process—changes in the flux fields of the magnetic domains on the disk surface, from positive to negative and back again, induce changes in the micro-current flowing through the head.

More information about how individual components of typical disk drives work is provided in the following glossary of terms.

*Actuator.* The mechanism that positions the head over the proper track on the drive. The class of actuators that is now most commonly used, because of its superior positioning ability, is called a "voice coil" motor. This operates on a

Downloaded from <https://www.cambridge.org/core>. IP address: 207.241.231.81, on 03 May 2019 at 03:42:24, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.2307/3116804>

principle similar to that used in telephones—an arm is moved in and out via electromagnetic forces. Voice coil motors have been made in linear and rotary designs, but the rotary design (which works like the arm on a phonograph) has become the dominant design because it requires less space. A much less expensive actuator mechanism is a *stepper motor*, in which a shaft rotates in discrete steps to new positions in response to changes in the surrounding magnetic field. Stepper motors were used primarily on low-capacity drives targeted to price-sensitive markets. *Torque motors* and *DC motors* were also used on a limited number of models in the low-to-moderate performance range.

*Areal Density.* The amount of information that can be stored in a square inch of disk surface, measured in megabits per square inch (mbpsi). This is determined by multiplying the number of bits of information storable along a linear inch of track (*bit density*) by the number of tracks per inch of disk radius (*track density*).

*Disk.* The round, rigid platter on which data is magnetically recorded. It is composed of a substrate, typically made of aluminum polished perfectly flat, coated with particles of magnetic metal oxide or thin metal films. These magnetic coatings are, in turn, coated with lubricating and protecting materials to prevent dislodging of information-bearing material out of the disk in the event that the head crashes into the disk surface.

*Drive.* The computer industry's term for the equipment that contains and rotates magnetic media—reels of tape, flexible (floppy) disks, or rigid disks—and that controls the flow of electronic information to and from those media.

*Embedded Servo System.* Mechanical shocks, differential thermal expansion, and a host of other factors can affect the accuracy with which an actuator can position a head over a particular track on a disk. Low-performance drives using stepper motor actuators got around this problem by spacing the tracks far enough apart that such subtle changes and misadjustments rarely caused the head to be mispositioned. High-performance drives, however, require a closed-loop feedback system to the actuator, so that the head can continuously be repositioned precisely over the proper track on the disk. This enables much greater track density. One way of keeping precise head-disk alignment was to dedicate one complete surface of one disk on the spindle to tracking information only. The head reading information off that track and feeding it back to the actuator motor provided such a closed-loop, continuous-adjustment mechanism. In an *embedded servo system*, track identification markers are written (embedded) on each individual track of each recording surface. This frees up for user information the entire surface that otherwise would have been reserved for tracking information only.

*Ferrite.* A magnetic compound composed of iron and oxygen. In disk drives, the primary use of ferrite has been as the core material around which fine copper wires were coiled to form an electromagnet in the head.

*Head.* A device that contains a tiny electromagnet, positioned on an arm extending over the rotating disk. When the direction of current through the head changes, its polarity switches. Because opposite magnetic poles attract, changes

in the polarity of the head causes an opposite change in the polarity of the magnetic material on the disk as it spins immediately beneath the head. The head writes information in binary code in this fashion. Heads read data in the opposite manner—changes in the magnetic flux field over the disk's surface as it spins beneath the head induce changes in the direction of current in the head, reversing the information flow. In rigid disk drives, heads are aerodynamically designed to fly a few millionths of an inch above the surface of the disk; they generally rest on its surface when the drive is at rest, take off as the disk begins spinning, and land when the disk stops again. Heads in floppy disk drives generally do not fly but glide on the disk's surface.

*Interface.* This term refers to the electronic circuitry through which the drive and computer communicate. A thorough description of the differences among interfaces is beyond the scope of this article. Originally, interfaces were custom-written by each drivemaker for each customer. Although some standard interfaces such as SMD emerged as 8-inch drives were used with minicomputers, the trend toward standardization was accelerated by Seagate Technology's ST412 interface, which required that the rate at which the drive took data off the disk was equal to the rate at which the drive could transfer data to the computer. Although low-cost and efficient, this system effectively put a ceiling on the bit density of the drive. Subsequent interfaces such as SCSI (used primarily with Apple computers), AT (used with IBM-compatible computers), and ESDI (used primarily with engineering workstations) decoupled these activities. With these interfaces, the drive could take data off the disk as rapidly as its designers wanted, cache it, and then transfer it to the computer as rapidly as the computer could accept it. This enabled much greater bit densities than had been possible under the ST412 interface. Other interfaces used on only a limited number of models were IPI-1, IPI-2, and ANSI.

*MFM.* An acronym for *modified frequency modulation*, an early coding technique used in writing data on disks, wherein a magnetic marker was placed on the disk to denote the beginning and ending of each individual piece of information.

*MIG Heads.* An acronym for *metal-in-gap*, a version of ferrite head wherein a strip of metal was deposited in the gap between the leading and trailing portions of the head. This strengthened the magnetic flux fields that could be created and sensed by the head, enabling data to be written and read on smaller domains on the disk surface.

*Oxide.* The term used in the industry for particles made from a compound of oxygen and a magnetic metal, such as iron, cobalt, and chromium. Oxide particles were used to coat mylar substrates to create magnetic tape and floppy disks and to coat aluminum disks in rigid or "hard" disk drives. The oxide particles are the media in which, through changes in the particles' magnetic polarity, data are stored magnetically. The particles generally have an elongated, needle-like shape.

*Photolithography.* The manufacturing process through which a desired pattern of one material is applied onto another substrate material. Typically, the substrate is first coated (by plating or *sputtering*) with the material from which the

final pattern is to be made. This is in turn coated with a light-sensitive monomeric material, called a *photoresist*. A mask of the desired pattern is then held over the photoresist, and the unmasked material is exposed to light, causing the exposed material to cure. The unexposed photoresist is then washed away. Through a subsequent series of etching and washing steps, only the desired material, in the desired pattern, is left on the substrate. Integrated circuits are built on silicon wafers and thin-film heads are built through photolithographic processes.

*PRML*. An acronym for *partial response, maximum likelihood*, a coding technique that has followed RLL and MFM recording codes.

*Recording Density*. See *areal density*.

*RLL*. An acronym for *run-length limited* recording codes, which enable data to be written more densely than was possible with MFM codes. Two versions of RLL codes have been used: 2,7 and 1,7.

*Spin Motor*. The electric motor that drives the rotation of the spindle on which the disks are mounted. In 14- and 8-inch drives, the spin motor was situated in the corner of the drive and drove the stack of disks via a pulley. In the 5.25-inch and subsequent drive architectures, a flat, "pancake" motor was developed and positioned beneath the spindle, whose rotation it drove directly.

*Spindle*. The shaft on which one or more disks was mounted.

*Stepper Motors*. See *Actuators*.

*Thin Film*. A continuous, very thin film (often only a few angstroms thick) of a material (often a metal) on another substrate material. This is generally applied through a process called *sputtering*, in which a substrate is placed at the bottom of a vacuum chamber. A target of the film material is then bombarded with electrons, which dislodge ions of the target material. These ions float like a vapor in the vacuum chamber and then gradually settle in a thin, continuous film on the surface of the substrate. This deposition technique is one of the early production steps in the manufacture of integrated circuits and thin-film heads. It is also the technique used to coat disks with very thin films of magnetic material.

*Torque Motors*. See *Actuators*.