Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

The Fruit Flies of Electronics

> When I began my search for an answer to the puzzle of why the best firms fail, a friend offered some sage advice. "Those who study genetics avoid studying humans," he noted. "Because new generations come along only every thirty years or so, it takes a long time to understand the cause and effect of any changes. Instead, they study fruit flies, because they are conceived, born, mature, and die all within a single day. If you want to understand why something happens in business, study the disk drive industry. Those companies are the closest things to fruit flies that the business world will ever see."
>
> —*The Innovator's Dilemma* by Clayton M. Christensen, pg. 3.

E*nthusiast*: "I believe applications of MEMS in electronic systems will foster a revolution that will transform the industry."

Skeptic: "What revolution? MEMS have been around for 50 years. Sure, there are MEMS accelerometers in airbag sensors, it's hardly a revolution; it took the automobile industry 30 years to introduce variable-speed windshield wipers."

To answer this skeptic, let's take the advice of Clay Christensen's friend and look at hard disks. Since the disk industry evolves rapidly, it pushes the limits of practical application in motors, actuators, sensors, bearings, lubricants, connectors, materials, circuit boards, electronics, packaging, and a myriad of other components. Are MEMS on the horizon for any important hard-disk components? Since the hard disk has been evolving and improving for 45 years (IBM introduced its RAMAC hard disk in 1956), it may be approaching fundamental barriers. If progress slows, what alternatives will challenge the hard disk's position as the most cost-effective digital mass storage medium? Will its successor be a microelectromechanical system?

What follows is not really about hard disks. You will see so much about hard disks that it will seem as if they are the subject of this report. They aren't. This month's newsletter is no more about hard disks than *The Innovator's Dilemma* is about hard disks. Clay Christensen wrote about how companies fail. I'm writing about how MEMS will invade electronic systems. Clay chose the hard disk industry because it evolves so rapidly that it is like studying genetics by observing the rapidly evolving generations of fruit flies. Choosing hard disks was an excellent idea. I'm choosing hard disks for the same reason: they are electronic fruit flies—they evolve so rapidly that we can use them to gauge how other electronic systems will evolve.

Hard disks seem boring. The hard disk is cheap. It gets faster and bigger every year. Ho hum. The "sweet spot" of the consumer hard-disk market is about $150. Access times have shrunk from 85 ms for the IBM PC/XT's 10 megabyte disk (1983) to 8.5 ms for IBM's 75 gigabyte Deskstar (2001). Disk capacity has been increasing 24 percent per year for *40 years*.

The hard disk is an electromechanical contraption with mechanical arms and magnetic sensors sweeping over the surfaces of a stack of spinning disks coated with magnetic material. The hard disk is a direct and seemingly inelegant way to store massive amounts of data. Surely, in the age of electronics, there are ways to store bits that don't involve motors, bearings, moving arms, electromechanical controls, air filters, and lubricants.

One would think so. And there has always been something on the horizon to displace the hard disk. Bubble memory would make the disk obsolete. Flash memory would do it. Then optical storage or holographic storage. It hasn't happened. How about an imminent challenge from Lilliput in the form of

In This Issue:

Vol. 1, No. 3
February 2001

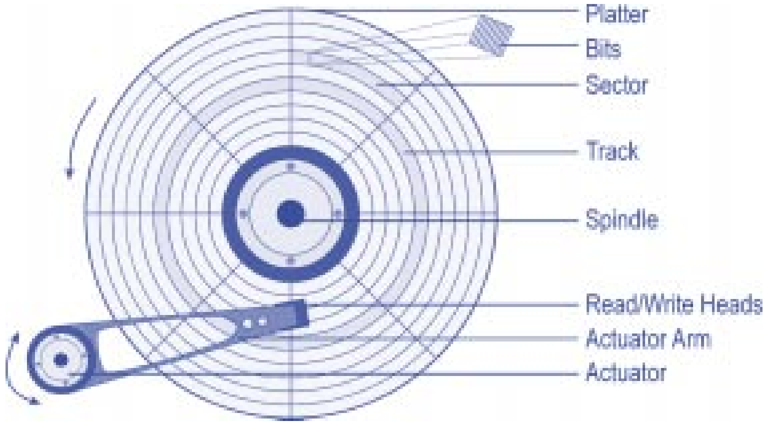

Figure 1. The hard-disk platter is divided into tracks and sectors. The sectors contain bits stored in magnetic particles embedded in the coating on the surface of the disk.

MEMS? MEMS will not displace the hard disk in the near term. In fact, MEMS will prolong the hard disk's reign by invading components inside the hard disk, helping increase its speed and capacity. Still, because of its rate of evolution, the hard disk is a great place to start an investigation of how microelectromechanical systems might first enter and then disrupt electronic systems.

The hard disk: how it works

Stick with me for a few paragraphs on how the disk works. Once we know something of how it works, we'll be able to project where it's going, how it's likely to get there, and when it's likely to encounter technical barriers to continued improvement. If the hard disk's rate of improvement slows as it approaches these barriers, its competitors may gain ground.

Figure 1 shows a hard-disk platter with its read/write actuator arm. The disk platter is a flat aluminum or glass plate coated with magnetic material on both sides. There's one actuator arm and its corresponding read/write head for each disk surface. The disk may have one to six (or more) platters (and twice that number of actuator arms and read/write heads). There's only one actuator motor, so all of the arms are stacked on top of each other and they all move together. A platter's surfaces are divided into tracks and sectors. Tracks are concentric rings of bits. A sector holds the number of bits that is the smallest amount read or written by the disk's electronics. As we will see, the figure isn't to scale. Also, I have shown eight sectors per track from the center of the platter to the perimeter. In modern hard disks, the platter may have more sectors on the outer tracks than it does on the inner tracks. A "bit" is a tiny part of a sector containing at least a few hundred particles of magnetic material. The read head detects the direction of alignment of the magnetic disturbance representing a bit in the material. For writing, a current in the write head forces the material into the magnetic alignment corresponding to a one or a zero as the region passes under the write head.

DynamicSilicon

Editors Nick Tredennick
Brion Shimamoto
Publisher Lauryn Franzoni
Associate Editor Keri Geiger
Designer Julie Ward
President Mark T. Ziebarth
Chairman George Gilder

DynamicSilicon is published monthly by Gilder Publishing, LLC. Editorial and Business address: P.O. Box 660, Monument Mills, Housatonic, MA 01236. Copyright 2001, Gilder Publishing, LLC. Editorial inquiries can be sent to: bozo@gildertech.com Single-issue price: $50. For subscription information, call 800.261.5309, e-mail us at help@gildertech.com, or visit our website at www.dynamicsilicon.com

The electronic systems market is primarily three overlapping segments based on major system design objectives (*Dynamic Silicon* January 2001): zero cost, zero delay, and zero power. The zero-cost segment is the consumer market; designers strive for minimum-cost systems. The zero-delay segment is the performance market; designers strive for minimum delay from request in to answer out. The zero-power segment is the mobile market; designers strive for minimum power dissipation. Hard disks are built for the zero-delay market. Different models address the desktop segment and the performance segment. For example, IBM's Deskstar drives serve the consumer desktop segment (the overlap of zero cost and zero delay) and its Ultrastar drives serve the performance segment (zero delay but not zero cost). Hard-disk manufacturers compete on speed and on storage capacity. The components of speed are access time and data transfer rate.

Access Time. Access time is seek time plus latency. Seek time is how long it takes to move the read/write head to the correct track, shown in Figure 1. Latency is the time it takes for the desired sector to rotate underneath the read/write head (on average, half a disk revolution). To reduce seek time, we can make the platter smaller (so the actuator arm doesn't have to travel as far) or we can make the actuator arm move faster by increasing the power of its motor, or by reducing the weight of the arm. To reduce latency, we can increase the speed of rotation. The left side of Figure 2 plots latency against rotation rate and shows how the latency decreases as the rotation rate increases. Slower hard disks serve consumer and desktop markets. Faster hard disks

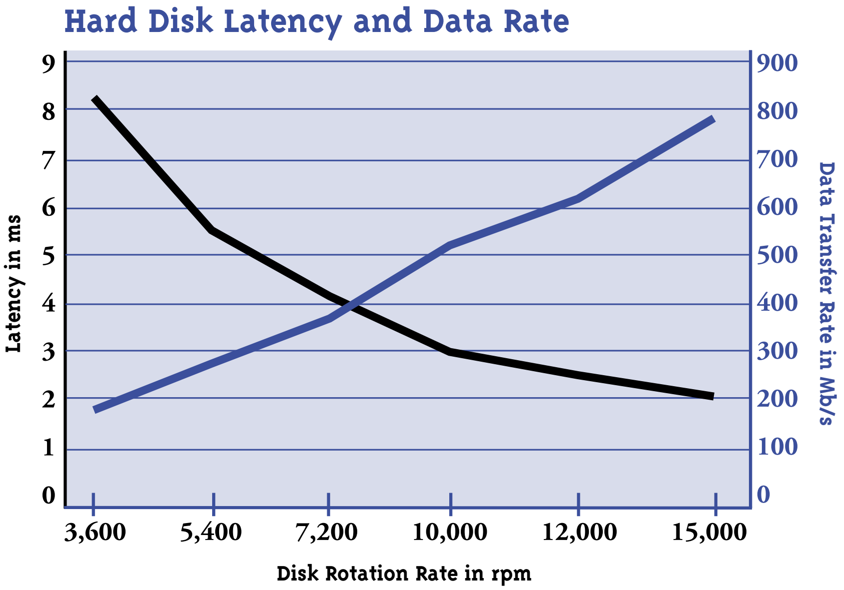

Figure 2. Increasing disk rotation rate raises the data transfer rate and decreases the latency. This figure is for a bit density of 500,000 bits per inch, written "500 Kb/in" in the track.

appear in servers and workstations. Seek times happen to be about 1.5 times the latency.

Data Transfer Rate. Data transfer rate is a measure of how fast the hard disk can move bits between the platter's magnetic particles and the disk's electronics. The rate is determined by the number of bits per inch in the track and by the rotation rate of the platter. The right side of Figure 2 plots data transfer rate against rotation rate, assuming a bit density of 500 Kb/in, this shows how data transfer rate increases as disk speed rises. The data transfer rate has grown rapidly because it is proportional to the *product* of the disk rotation rate and the number of bits stored on a track, *both* of which are increasing over time. Doubling the rotation rate and doubling the number of bits per track increases the data transfer rate by a factor of four. The number of bits on a track has risen rapidly with improvements in platter coatings and with improvements in read/write heads.

As the data transfer rate increases, data transfer protocols between the hard disk and the computer must improve. The Ultra ATA/66 data transfer protocol, for example, popular at the high end of the desktop market this year, is capable of 528 million bits per second (Mb/s). At the bit density illustrated in Figure 2 (500 Kb/in) the data transfer rate for a 12,000- or a 15,000-rpm disk will saturate an Ultra ATA/66 interface. This year's 15,000-rpm hard disks come with faster interfaces popular in workstations and servers. Hard-disk data transfer rates will soon outrun the server's interface protocols.

Storage Capacity. We get more capacity by making the bits smaller, not by making the container bigger. Smaller bits mean more bits per track and more tracks per platter, which adds up to more storage capacity or, as the industry calls it, more "areal density." In just the last 10 years, linear density has grown from about 40 Kb/in to about 500 Kb/in, while track density has grown from about 2 K tracks/in to about 40 K tracks/in. Areal density has grown from less than 0.1 gigabits per square inch (a billion bits per square inch) to more than 20 Gb/in^2 by 2000. Areal density could reach 50 Gb/in^2 this year.

You make bits small (and increase density) either by shortening the bit or by narrowing the track. In 1990, as Figure 4 shows, the track width was 20 times the length of a bit. In 2001, the track width will be about 10 times the bit length, so the width is shrinking faster than the bit length, but it has a long way to go. Track width could be improved to be about four times the bit length for longitudinal recording (the kind we have been using for 45 years). Track width and bit length are sensitive to progress in different areas.

Read/write head improvements lead to shorter bit lengths. Passive MEMS (MEMS that don't have moving parts), in the form of the thin-film read/write head, invaded the hard disk in the early '80s. In the early '90s, IBM invented an improved passive MEMS thin-film write head

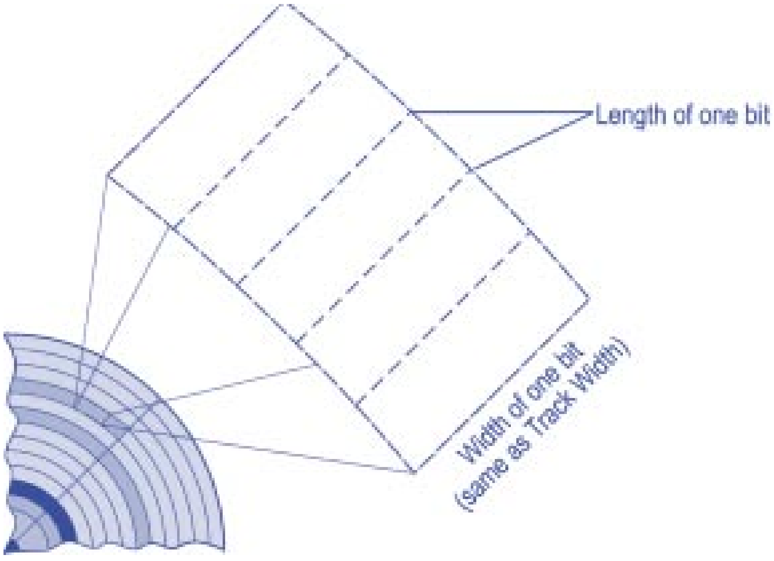

Figure 3. Illustration of bit length and track width.

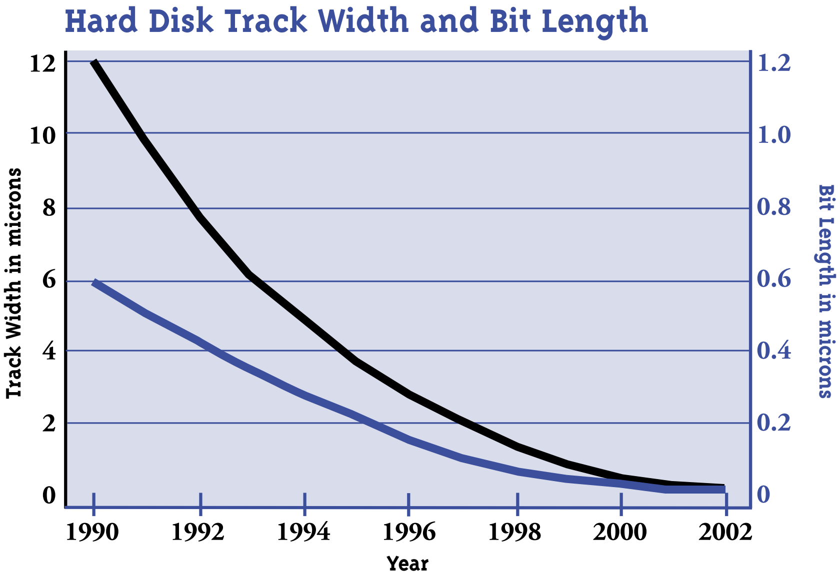

Figure 4. In 1990, bits on the hard disk were about 20 times as wide ("track width") as they were long ("bit length"). In 2001, track width is still about 10 times bit length. A micron is one thousandth of a millimeter.

coupled to an MR (magnetoresistive) read head. IBM introduced the first MR head in 1991 in its Corsair 8-platter 3.5" 1 GB hard disk. In 1997, IBM began the transition to GMR (giant magnetoresistive) read heads in the Titan 16.8 GB hard disk. These improvements in read/write heads let the bits get smaller, improving areal density. Detection also improves as the disk spins faster and as the head moves closer to the disk's surface. Today's read/write heads fly at 0.01 microns above the disk's surface (about *a tenth of* the line width in today's integrated circuits and 5,000 times smaller than the diameter of a human hair). This shows that mechanical systems can operate at peer tolerances with semiconductors. Future generations will bring this height down to the thickness of the surface lubricant's molecule.

If bit length drops by improving the sensitivity of the read/write head, the track width improves by positioning the read/write head more accurately. The hard-disk controller must slew the actuator arm and its read/write head to the desired track and hold it there. The controller wants to get the read/write head in position in the minimum possible time. (That's seek time.) It's something like swinging a cue stick by its fat end to point to something with the tip—the controller is a long way from the read/write head. Once the read/write head is positioned over the track, it has to be held close to the center of the track despite electronic noise and physical vibration. The need for rapid, precise positioning and the magnitude of the noise and vibration in the system make reducing the track width particularly difficult. That's why the track width is still 10 times the bit length.

MEMS to the rescue

MEMS invaded the read/write head to reduce the bit length. MEMS will soon invade the actuator system to reduce the track width.

Rotational vibrations cause alignment errors between the track and the read/write head, limiting the track width. STMicroelectronics is attacking the rotational vibration problem with MEMS. Its L6670 capacitive rotational accelerometer, placed anywhere in the drive's case, senses vibrations (by measuring minute rotations). Using the L6670's output, firmware computes a corrective signal for the motor to control the read/write head position. Better tracking enables narrower tracks. This is only the beginning for STMicro, which has created a business unit and has dedicated a manufacturing line to support MEMS applications. Its first products, the microaccelerometer and a head-positioning microactuator, another device to reduce tracking error, are for the hard-disk market, but it also expects to produce MEMS for optical markets and for wireless RF markets.

In today's hard disk, the read/write head is fixed to the end of the actuator arm. To keep the head positioned in the middle of the track, the actuator motor moves the actuator arm. The head-positioning microactuator is a small linear-displacement motor on the read/write head end of the actuator arm. The actuator motor points the actuator arm in the general direction of the middle of the track. The microactuator, using electrostatic forces acting between silicon plates, moves the read/write head to the middle of the track and keeps it there. This is like swinging your arm in the general direction of something and using your finger to point precisely. The microactuator can move the read/write head across several tracks to either side of the target track, opening the possibility for greatly reduced seek time for accesses within its range. Last April, Texas Instruments funded MEMS research at the University of California, Irvine to adapt a TI-developed silicon micromotor for use as a head-positioning microactuator. Other universities, including University of Tokyo, University of California, Berkeley, and Caltech are also working on head-positioning microactuators. In addition to TI, Applied Magnetics, HP, IBM, Magnacomp, Maxtor, Quantum, Read-Rite, and Seagate also sponsor university MEMS research.

Actuator arms are typically stamped stainless steel. Researchers at U.C. Berkeley and at IBM's Almaden Research Center are experimenting with structures for silicon actuators. The actuator arm would be better in silicon. In a silicon structure known as HexSil, top and bottom flat sheets cover a honeycomb of hexagonal cells to construct a strong, light actuator arm. It has nearly the same strength as stainless steel, with 1/40th of the mass. Less mass means less inertia and less momentum. A lighter arm accelerates quicker, moves faster, and stops quicker. This reduces the seek time and reduces tracking error, enabling more tracks per inch. In addition, since the fabrication of such actuator arms is compatible with standard semiconductor processes, other components of the assembly could be fabricated into the structure, reducing assembly cost and improving reliability.

Researchers at the University of Tokyo are experimenting with metal alloys called "shape metal alloys." These alloys change shape when a current is passed through them. Tiny shape-metal-alloy legs placed near the read/write head's surface on the actuator arm, prevent the surface from touching the disk surface when the disk is stopped. This enables the use of a smoother disk surface and a lower read/write head flying height. Lower flying height leads to reduced bit length.

MEMS will improve the hard disk's performance, storage capacity, and cost. But they cannot save it from its fundamental barriers.

The superparamagnetic limit

The platter's glass or aluminum substrate is coated with a material containing magnetic particles. Each bit on the platter contains at least a several hundred magnetic particles (about 500 to 1,000) because the bit's signal to noise ratio depends on the number of particles in a bit (more particles is better). The bit's area is the track width times the bit length. Each particle is a little magnet with its magnetic axis approximately aligned with the track (circumference). The write head magnetizes particles in one direction for a one, and in the opposite direction for a zero, as seen in Figure 5. Since the magnetic material is a uniform coating, the bit area shares boundaries with two other bits and with the adjacent tracks. This uncertain border region is one reason the bit area must contain several hundred magnetic particles. As the bit area shrinks, the particle size must decrease to maintain the particle count.

The superparamagnetic limit is reached when the particle size is small enough to make thermally induced magnetic reversals, and thus memory errors statistically likely. The describing equation is an exponential function of temperature and of particle volume. The superparamagnetic limit is *very* sensitive to particle size. Reducing the particle volume by half could change the average time between random reversals from years to nanoseconds. If the particles in a bit experience random reversals (errors) every few nanoseconds, the bit is scrambled soon after it is written. The material doesn't retain information. Shrinking the bit area stores more gigabits per square inch, but it also shrinks the particle size and it brings us closer to the superparamagnetic limit.

IBM researchers John Best and David Thompson, in an article in the May 2000 *IBM Journal of Research and Development*, estimate that the superparamagnetic limit can be pushed to 100-200 Gb/in^2. Last year, manufacturers shipped hard disks with areal densities of about 20 Gb/in^2. In 2002, areal density will reach 80 Gb/in^2. So it will only be a few years until the superparamagnetic limit slows progress in the hard disk's speed and storage capacity.

Several alternatives are there for pushing back the projected limit (it's been predicted and pushed back before). One alternative is to pattern material on the platter to isolate bits from their neighbors, but that's more expensive than simply coating the platter uniformly. Another possibility is to convert from longitudinal encoding to vertical encoding (standing the particles vertically in the coating), but this would entail redesign of the read/write heads. A third possibility is to find magnetic materials that will be less susceptible to random reversals. There are others.

The alternatives for continued improvement in hard disks are *all* difficult and everyone knows it. Incremental improvements won't be enough. That's causing renewed interest in alternatives. In the near term, MEMS will invade the hard disk as a result of ordinary manufacturing progress. MEMS cannot save the hard disk from the superparamagnetic limit, however. In the long term, new mass storage will

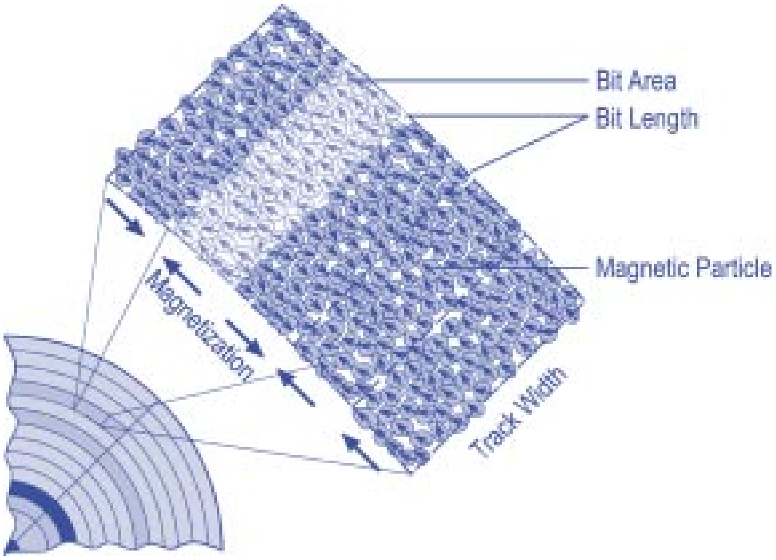

Figure 5. A few hundred magnetic particles constitute a bit on the disk's surface coating.

overtake the hard disk. But usurpers will have a difficult time reaching the hard disk's cost effectiveness. For consumers, hard-disk storage is below a third of a cent per megabyte. At that rate, it's cheaper than paper or film. So to get started, to get the first products to market and start moving along the learning curve, new mass storage will have to begin in a market segment that today's hard disk does not service well, such as small portable devices.

MEMS-based data storage

According to Richard Carley, Gregory Granger, and David Nagle of Carnegie Mellon University, "MEMS-based storage systems are potentially a whole new storage technology capable of a dramatic decrease in entry cost, access time, volume, mass, power dissipation, failure rate, and shock sensitivity. More important, these devices can integrate computation with storage, creating complete system-on-a-chip solutions—including mass storage." (See "MEMS-Based Integrated Circuit Mass-Storage Systems," *Communications of the ACM*, Nov. 2000, pp. 73-80.) I'm not sure that entry cost will be dramatically lower, but the rest of it is logical. MEMS can offer chip-scale mass storage. A MEMS-based mass storage system would certainly have low access time, low volume, low mass, and low power dissipation relative to its hard-disk cousin. Under the sponsorship of DARPA (Defense Advanced Research Projects Agency) and with aid from IBM and Intel, the paper's authors are working on just such a system.

Carnegie Mellon University's (CMU's) MEMS-based data storage. Imagine two chips each a little smaller than a postage stamp, positioned one above the other. The lower chip is an array of read/write heads. The upper chip contains an array of magnetic storage bits. This upper chip moves around under the control of miniature X and Y drive motors. In this way, the stored bits are scanned across the read/write heads in the X and Y directions. The whole chip, including the X and Y drive motors is just 10 millimeters on a side. The stored bits in the upper chip are arranged in a checkerboard pattern, with one read/write head for each square. As the array is scanned in the X and Y directions, each read/write head can read or write bits in its square. The whole two-chip mechanism, including the drive motors, is fabricated using semiconductor processes. Because the distances are small and because there's a whole array of read/write heads, the seek time (the time for a read/write head to be in the right place) for this MEMS-based storage device is about one tenth that of a hard disk (~0.6 ms versus ~10 ms). Because there are many read/write heads that can be transferring data at once, the aggregate data transfer rate can be high. CMU's proposed system has 6,400 read/write heads, of which 1,280 can be simultaneously active. Since the bits are stored in magnetic particles similar to those used in a hard disk, it is subject to the same superparamagnetic limit, which constrains the bit area. This prototype system stores 4.0 GB.

The IBM "Millipede." IBM's Millipede, like the CMU system, uses two small stacked chips. Unlike the CMU system, the IBM Millipede does not use magnetic bits for storage, and is therefore not subject to the superparamagnetic limit. Instead, it uses principles of the atomic force microscope (AFM) in an array of sharp (millionths of a millimeter-scale) probe tips. This is an atom-scale technology. Since the probes and mechanisms of the storage system are so small, the bit storage areas can be atom scale. AFM probe tips thermomechanically write and read indentations in nanometer-thick polymer film coating a silicon substrate. Since AFM probe tips operate on a microsecond scale, a thousand times slower than the read/write heads of today's hard disks, researchers aggregate probes to achieve faster data rates. IBM researchers believe that the principles used in the Millipede could achieve a data storage density of 500 Gb/in^2. Devices the size of integrated circuits could store gigabytes. IBM has demonstrated a 32 x 32 array chip. But, it's IBM research, not a product division, so there's no schedule for product implementation.

HP's Atomic Resolution Storage. HP's atomic resolution storage (ARS) may be even more ambitious than IBM's Millipede. HP's research project may be 10 years from practical devices. Like the IBM project, it uses atom-scale storage. Its storage material will have two stable phases at room temperature, crystalline and amorphous. An array of atom-scale probe tips is suspended about a micron above the storage surface and connected to scanning motors. A probe tip creates a one or a zero by a thermally controlled phase change of a spot with a directed beam of electrons. HP hopes to achieve storage densities greater than 1,000 Gb/in^2 (a terabit).

Lessons from the hard disk

The hard disk progressed from the 24-inch platter in 1956, to the 1-inch platter in 1999. MEMS brought micro-scale production to the hard disk in 1979 with the introduction of the thin-film head. Mechanical constraints dominate hard-disk design. Increasing speed and storage capacity requires decreases in the size of platters, actuators, sensors, and other components. MEMS permit parts that are small enough to be batch fabricated with semiconductor processes. The move from large, wire-wound heads to thin-film heads, for example, brought semiconductor manufacturing efficiencies to the disk's read/write heads. Other components will soon follow.

Stamped metal components, such as the disk's actuator arm, may be cheap, but require assembly to attach electronic components. Silicon components are lighter than their steel counterparts, improving performance. Silicon components may be cheaper because they can be batch fabricated, in efficient shapes unavailable to stamped and folded parts. Silicon components also offer the possibility of integrating electronics components to reduce assembly costs.

The transition from molded, stamped, and folded components to silicon components batch fabricated in a semiconductor process opens the door for continued miniaturization.

The hard disk is the interface between the computer's bits and bytes and the analog world of magnetic particles on the disk's surface. For the first two decades of hard-disk development, we bridged that gulf with macro-scale components that grew continually smaller. With the introduction of the thin-film head by IBM in 1979, the hard disk began a transition to micro-scale components. Micro-scale components will continue to displace other macro-scale components in the hard disk. They bring with them a new range of miniaturization.

MEMS aren't invading the hard disk because MEMS are cool; MEMS are invading because they make economic sense. The hard disk and its components are shrinking, driven by the desire for speed and storage capacity—objectives that can be achieved by decreasing size and increasing precision. Hand-made components gave way to machine manufacturing, which offered miniaturization, precision, and the economy of mass production. Silicon offers miniaturization and precision beyond what could be achieved with macro-scale manufacturing processes and retains the production economy of batch fabrication. In addition to these properties, silicon offers increased strength and reduced weight. With silicon, electronic and structural components can be integrated to avoid expensive assembly costs and to improve reliability. These are advantages whose applications extend beyond the hard disk

The hard disk is converting to micro-scale production as it incorporates MEMS components. Other electronic systems will soon follow. In the hard disk, MEMS are at the interface between the computer's bits and bytes and the analog world of tracking position and magnetic particles.

As the world divides into tethered and untethered (mobile) devices, mobile devices will be the collectors and consumers of data. MEMS are consistent with the zero-cost, zero-delay, and zero-power vision of untethered devices. Mobile devices will emphasize power conservation and portability. MEMS provide the transition from macro-scale sensors to miniature batch-fabricated sensors. MEMS will invade mobile systems as they have the hard disk. MEMS in mobile systems convey the advantages of batch fabrication, component and electronics integration, reduced weight, improved efficiency, and progressive size reduction.

MEMS will soon invade cell phones, personal digital assistants, GPS receivers, MP3 players, and electronic toys where they will convey to these applications the same advantages they have in the hard disk. Critical for mobile devices, MEMS help lower power consumption because their small scale results in economy of interaction with the real world. MEMS will help make portable a whole host of applications that today are confined to benches and laboratories.

The cell phone is ripe for a MEMS invasion. A cell phone contains 300 to 400 discrete components. Each of these separately fabricated components contributes to the variability among units and each contributes to difficulties in reliability and in manufacturing. Many of these components could be consolidated onto MEMS chips. Consolidated components contribute to reliability and to manufacturability. Further, component variation is controlled for components on the same chip.

Once confined to laboratories, MEMS will bring portability to biomedical/chemical analysis. In analysis of liquids and gases, for example, MEMS can work with minute samples. Because the on-chip analysis lab is batch fabricated, it can be replicated at little incremental cost. This leads to quick, accurate analysis of even small, dilute samples. It will lead to a huge change in the economies of lab tests. A test might use a disposable chip rather than an analysis lab and a technician.

Think of MEMS as a step in the industrial revolution that takes us from macro-scale component production to micro-scale component production. The move to micro-scale production requires investment in developing tools, designs, and methods. The payoff for that investment is in capabilities inaccessible to macro-scale production and in a continued range of reduced sizes and improved efficiency. ICs let us shrink electronics. MEMS will let us shrink their mechanical equivelants in the same way.

Brion N. Shimamoto

Nick Tredennick and Brion Shimamoto
15 February 2001

This month the *Dynamic Silicon*'s Company List has made the change from fractions to decimals.

Dynamic Silicon Companies

The world will split into the tethered fibersphere (computing, access ports, data transport, and storage) and the mobile devices that collect and consume data. Dynamic logic and MEMS will emerge as important application enablers to mobile devices and to devices plugged into the power grid. We add to this list those companies whose products best position them for growth in the environment of our projections. We do not consider the financial position of the company in the market. Since dynamic logic and MEMS are just emerging, several companies on this list may be startups. We will have much to say about these companies in future issues.

Altera (ALTR http://www.altera.com)

Altera and Xilinx together dominate the programmable logic business, with almost seventy percent of the CMOS PLD market. Both companies are aggressive and competitive. Sixty-six percent of Altera's revenue comes from the rapidly growing communications segment (Telecosm companies) and an additional sixteen percent comes from the electronic data processing (EDP) segment. Altera is positioned to be a major supplier in tethered applications such as the base stations that support the mobile devices.

Analog Devices (ADI http://www.analog.com)

Analog Devices is a leader in analog electronics for wireless RF and communication, MEMS for automotive applications (accelerometers, pressure sensors, transducers), and in DSPs.

ARC Cores (ARK (London) http://www.arccores.com)

ARC Cores makes configurable processor cores. Configurable processors allow the application engineer to adapt the processor's instruction set to the requirements of the problem. Conventional microprocessors have fixed instruction sets.

Cypress (CY http://www.cypress.com)

Cypress Microsystems builds components for dynamic logic applications. Cypress also builds MEMS and is a foundry for MEMS.

QuickSilver Technology, Inc. (* http://www.qstech.com)*

QuickSilver has the potential to dominate the world of dynamic logic for mobile devices (untethered). While many companies work on programmable logic and on "reconfigurable computing" for tethered applications, QuickSilver builds adaptive silicon for low power mobile devices.

SiRF (* http://www.SiRF.com)*

SiRF builds RF GPS chips for the mobile market. It is a world leader in development of integrated GPS receivers.

Transmeta (TMTA http://www.transmeta.com)

Transmeta makes new generation microprocessors that use closed-loop control to adapt to problem conditions in an x86-compatible environment. This enables Transmeta's microprocessors to save power over conventional microprocessors from AMD and Intel. The base instruction set is not available to the application engineer.

Technology Leadership	Company (Symbol)	Reference Date	Reference Price	Current Price	52-Week Range	Market Cap.
General Programmable Logic Devices (PLDs)	Altera (ALTR)	12/29/00	26.31	30.00	19.62 - 67.12	12.5B
Dynamic Logic for Mobile Devices	QuickSilver Technology, Inc. (none*)	12/29/00				
MEMS Foundry, Dynamic Logic	Cypress (CY)	12/29/00	19.69	24.26	18.25 - 58	3.2B
RF Analog Devices, MEMS, DSPs	Analog Devices (ADI)	12/29/00	51.19	50.96	42.63 - 103	18.5B
Configurable Microprocessors	ARC Cores (ARK**)	12/29/00	£3.34	£2.36	£2.01 - 4.58	£791.6M
Silicon for Wireless RF, GPS	SiRF (none*)	12/29/00				
Microprocessor Instruction Sets	Transmeta (TMTA)	12/29/00	23.50	32.25	17 - 50.88	4.1B

*QuickSilver and SiRF are pre-IPO startup companies.

** ARK is currently traded on the London Stock Exchange

NOTE: This list of Dynamic Silicon companies is not a model portfolio. It is a list of technologies in the Dynamic Silicon paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. The authors and other Gilder Publishing, LLC staff may hold positions in some or all of the companies listed or discussed in the issue.