Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

MEMS-based Storage

Wasn't "The Fruit Flies of Electronics" (February 2001 *Dynamic Silicon*) about hard disks? Is that two issues in three months about storage? The answer to *both* questions is "yes and no." "Fruit Flies" wasn't really about hard disks; it was about how MEMS will invade electronic systems. Hard disks were the example because hard disks evolve so rapidly. It was to show how MEMS are invading the hard disk and then, by extension, how they will invade systems that evolve slower. This issue *is* about MEMS-based storage, but not primarily because there are huge markets looming for its products. Rather, seeing the developments in MEMS-based storage, we will learn lessons that apply to other MEMS applications.

Two forces encourage the development of MEMS-based data storage: proliferating portable devices and the superparamagnetic limit. Small devices need cheap, rugged, capacious data storage. Hard disks are within a few generations of reaching the superparamagnetic limit, where bit sizes are too small to remember their data (thermal noise randomizes the bits). The search has begun for alternatives to longitudinal magnetic recording.

The world is splitting into tethered and untethered devices. The fibersphere, which caters to tethered devices (ones plugged into the power grid), will be the repository for computing, access ports, data transport, and storage. Untethered devices will be the collectors and consumers of data. Leading-edge untethered devices lie in the overlap of zero cost (consumer market), zero power (long battery life), and zero delay (compute-intensive tasks). Tethered and untethered devices need storage. Since 1956, the industry's data storage device has been the hard disk. A look at the best available hard disks for portable devices shows why hard disks are not ideal.

For portable devices, such as cameras, PDAs (personal digital assistants), and MP3 players, today's leading-edge hard disk is IBM's 1-GB Microdrive. The IBM Microdrive's one-inch hard disk spins at 3,600 rpm, giving it an average latency of 8.33 ms (milliseconds). Its average seek time is 12 ms. It has a linear recording density of 435 kb/in and a track density of 35,000 tracks/in, giving it an areal density of 15.2 Gb/in^2. The superparamagnetic limit, which is the maximum density for stable longitudinal magnetic storage, is estimated to be 100-200 Gb/in^2, so the Microdrive could grow to 6-12 GB. The Microdrive burns 66 mW (milliwatts) in standby and 825 mW when it is writing. According to Pricewatch (www.pricewatch.com), I can get a 1-GB IBM Microdrive for less than $400. This is 3 to 4 times the price of a 40-GB desktop hard disk, making the price per-gigabyte premium 120 to 160 times. The storage capacity premium in a portable device is forty times.

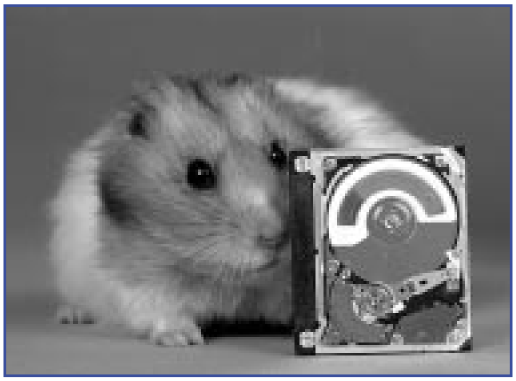

Fig. 1. IBM's one-inch, one-gigabyte Microdrive.
Courtesy of International Business Machines Corporation. Unauthorized use not permitted.

IBM's Travelstar 48GH is a good example of a leading-edge hard disk for laptop computers. It is a 48-GB 2.5-inch hard disk built to conserve power. The drive, which spins at 5,400 rpm, has a series of operating modes beyond simple reading and writing. These modes, in order of decreasing power use, are startup, write, read, performance idle, active idle, low power idle, standby, and sleep. Power dissipation in these modes ranges from 5 watts for startup to 100 mW for sleep. The stages of idle, standby, and sleep use less power, but they trade power for time: it takes seconds to get from sleep to a read or write. The disk's average latency is 5.5 ms, its average seek time is 12 ms, and its maximum data trans-

In This Issue:

Vol. 1, No. 5
May 2001

fer rate is 241 Mb/s. Pricewatch couldn't find the 48-GB Travelstar, but the 30-GB version is about $300.

By contrast with hard disks built for portable applications, IBM's Ultrastar 36Z15 is built for performance. It spins at 15,000 rpm. Average latency is just 2 ms, average seek time is 4.1 ms, and its maximum data transfer rate is 453 Mb/s. It burns 35 W on startup and 13.5 W in idle. These specs highlight the performance lag in hard disks for portables (whose emphasis is power conservation).

Hard disks use too much power. They conserve power by going to sleep and then they take a long time—seconds—to wake up and get to work. Hard disks for portable devices trail desktop and server units in performance (it takes power to spin the disk fast and to move the read/write heads rapidly).

Perhaps Flash memory is the answer. Its density will grow and its cost will decrease with Moore's law. It doesn't use power when it's not active. It has no moving parts to wear out. It sounds good, but the per-gigabyte price for Flash memory is currently 400 to 600 times the price of desktop hard disks. The storage density isn't good either. Intel claims to have the industry's smallest Flash cell. In a 0.18-micron process, Intel's Flash cell is 0.32 microns^2. If cell size scales directly with process progress (an optimistic assumption), it will be 0.17 microns^2 in a 0.13-micron process (planned for introduction in 2002) and 0.10 microns^2 in a 0.10-micron process (not yet on Intel's public roadmap for Flash). Even at 0.10 microns^2, it is equivalent to only 6.4 Gb/in^2—far below the density of hard disks.

Flash memory and miniature hard disks are too expensive and they are too small in absolute storage capacity to meet the needs of portable devices such as cell phones, PDAs, and MP3 players. Perhaps MEMS can combine the dense storage, fast data-transfer, and low cost of hard disks with the ruggedness, small size, fast access, and low power of semiconductors.

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: 291A Main Street, Great Barrington, MA 01230. Copyright 2001, Gilder Publishing, LLC. Editorial inquiries can be sent to: bozo@gildertech.com Single-issue price: $50. For subscription information, call 800.229.2573, e-mail us at dynamic@gildertech.com, or visit our website at www.dynamicsilicon.com

MEMS-based storage

In today's hard disk, a magnetic disk spins under a read/write head that arcs across concentric data tracks on the disk's surface. Since MEMS can have moving parts, could we duplicate the mechanisms in today's hard disk and bring to it the advantages of batch fabrication and miniaturization? Probably, but it wouldn't be the best way to build MEMS-based storage. Bearings are not easy to make and there is a problem with "stiction" (small moving components want to stick to each other). More importantly, if we are changing the "means" of production—mechanical assembly to semiconductor processing—we should look for ways to take advantage of the new process. For example, today's hard disk has one expensive read/write head per storage surface. With a semiconductor process, it won't be substantially more expensive to fabricate an array of 1,024 or 4,096 read/write heads than it is to build only one. Thus, instead of rotating the media and moving a read/write head, we might keep the array of read/write heads stationary and move only the media in a way that effects both the "track" positioning and the read/write motion.

Two chips. MEMS-based storage is two chips. Figures 2 and 3 show one way of doing it. The upper chip is an array of read/write heads. The lower chip is the media (storage array) and the actuators to move it. In the example of fig. 2, I illustrate electrostatic actuators. Four comb-like actuators bracket the media in the X and Y directions. Fingers from the media extend between the fingers of the actuators, increasing the area for electric charge interaction. The media is suspended in a way that allows movement in the X and Y directions. The media, which is grounded (zero volts), moves in response to voltage applied to the actuator combs. As the media moves, each read/write head scans a rectangular area where its bits are written.

Media. We have to decide how to store bits. Hard disks use magnetic storage. That's an option for MEMS-based storage, but, for a new design, we should consider alternatives. Alternatives sidestep the fast-approaching superparamagnetic limit and may offer density beyond that achievable with magnetic storage. For example, consider a phase-change material, which represents ones and zeroes as crystalline and amorphous phases of the material. Hewlett-Packard's experiments with phase-change media are expected to achieve 1,000 Gb/in^2. For really dense bit packing, consider removing individual atoms; vacant spots in a coating could represent zeroes.

The media for storing bits limits the choices for how to read and write the data. A phase-change media, requires heating to induce the phase change, plus some means for detecting the difference in material phase (e.g., differences in electrical resistance or in optical properties between the amorphous and crystalline phases).

Read/write array. We have to decide on the number of read/write heads and on how many can be active at a time. With 4,096 read/write heads simultaneously active, we get high data transfer rates, but then we need 4,096 input/output pins for the array and a Honda generator to supply the current. With fewer than 4,096 read/write heads active at a time, we have to decide how to share their input/output pins and how to organize the data for efficient access.

Motion. We have to decide how to move the media. We could move the media with electrostatic-comb actuators illustrated in fig. 2 or we could move the media magnetically. Moving it electrostatically restricts the distance that the media can move more than moving it electromagnetically. Electromagnetic actuation requires continuous currents and may need bulky magnets. Part of deciding how to move the media involves decisions about how to organize the data. What, for example, is a "track" and what kind of scanning pattern is best?

Since there are plenty of companies and universities that have worked on MEMS-based storage, let's look at the choices some have made.

Carnegie Mellon University (CMU). Researchers at CMU are building prototypes of a MEMS-based storage system. The CMU vision is a module with the area of a postage stamp that contains the CPU, RAM, MEMS-based storage, and the means to communicate. The cost goal for the module, which would contain about 10 GB of data storage, is $10 to $30 by about 2005. CMU researchers plan to cut power dissipation by a factor of one hundred and to improve average access times by a factor of ten over conventional hard disks. Achieving these goals would certainly encourage the proliferation of small, capable portable devices.

CMU's media uses perpendicular magnetic recording. This is similar to the longitudinal recording employed by today's hard disks. The difference is that in perpendicular recording the bits stand on end. This enables denser packing. But standing the bits on end forces redesign of the read/write head. And, since the device uses magnetic storage, the superparamagnetic limit will determine maximum storage density.

The stationary array has 6,400 read/write heads of which 1,280 can be simultaneously active. Each read/write head has a data transfer rate of 700 Kb/s, giving the storage array an aggregate data rate of 896 Mb/s. In magnetic storage, read/write heads do not contact the media surface. Due to variations in media surface height and to variations in the individual read/write heads, the CMU system uses active vertical positioning. Each read/write head is on the end of a cantilever beam that is electrostatically positioned above the media.

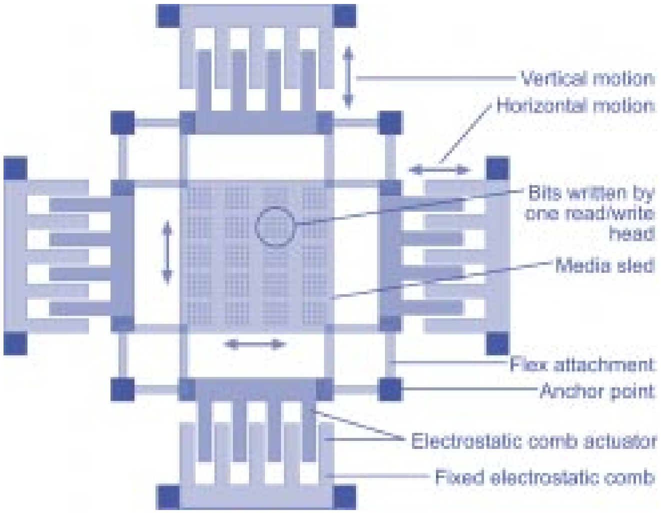

Fig. 2. The movable lower chip holds the magnetic media and the actuators. The stationary upper chip (not shown) holds the read/write heads.

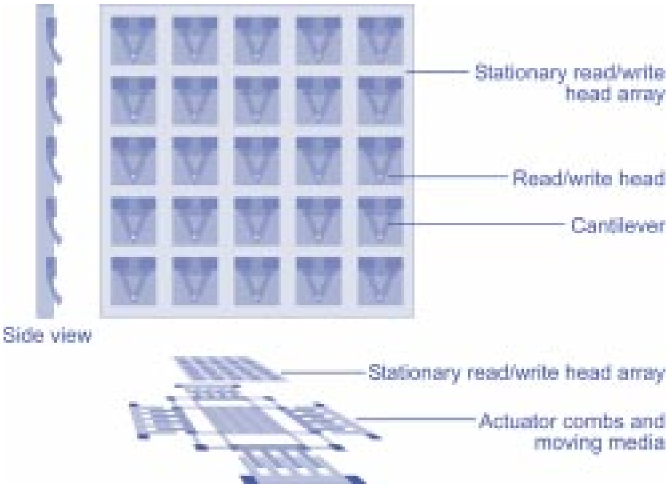

Fig. 3. Each read/write head accesses its own area of the media.

Electrostatic-comb actuators move the media in the horizontal plane. Electrostatic charge differences between the comb fingers and the media move the media in the X and Y directions.

IBM. Engineers at IBM's Zurich research center have built MEMS-based storage prototypes. The media for IBM's "Millipede" is a thin polymer coating on a silicon substrate. Storage is thermomechanical. The read/write head, which is in contact with the polymer coating, melts a dent in the surface to create a bit. The read/write head measures thermal conductance to determine whether it is in a dent (heat transfer is greater if the head is in a depression). Since the mechanism is not magnetic, it is not subject to the superparamagnetic limit. IBM researchers expect to reach densities of 500 Gb/in^2.

The stationary array has 1,024 read/write heads. Individual heads can write a bit and can read bits, but cannot erase individual bits (dents). Melting erases all the bits in an area. Since the read/write heads are in contact with the media, there's no need for individual height adjustment. Mechanical integrity and wear are potential difficulties, however, since the read/write heads contact the media. Data rates for each read/write head are a few megabits per second for reading and 100 Kb/s for writing.

Electromagnetic actuators move the media in the horizontal plane.

Samsung. Samsung projects the market for MEMS-based storage to be $3 billion by 2005 (about the same as the 2005 market for compact Flash memory). The Samsung Advanced Institute of Technology is designing a module for MEMS-based storage and, like IBM, has built at least one prototype. For its storage module, Samsung adds two signal processing chips to the media and read/write head array chips. Samsung plans to package four of these storage modules, together with control chips, on a SIMM (single in-line memory module). In the prototype, each module is 1 GB, giving the SIMM a 4 GB capacity. The media for Samsung's device is a coating of non-volatile ferroelectric material. Ferroelectric materials exhibit electrical polarization, which can be induced by a current and detected piezoelectrically. Samsung's engineers expect to reach storage densities of several hundred gigabits per square inch, with a terabit (1,000 Gb) per square inch as the ultimate goal.

The stationary array has some 2,000 read/write heads. Data transfer rate is listed (along with other module characteristics) as greater than 1 Mb/s, which seems about right per read/write head, but would be slow for the module.

Electromagnetic actuators move the media in the horizontal plane. The unit uses permanent magnets to stabilize the media. Power dissipation for the prototype module is one watt, with the goal to improve that to "several tens of milliwatts."

Nanochip. The information about Nanochip Corp. comes from press releases, from the web site (www.nanochip.com), and from U.S. Patent #5,453,970 (filed in 1993). Nanochip wouldn't talk to me when I called. "We're about a month from disclosing something." I don't think it would have mattered whether I called six months ago or six months from now, the answer would have been the same. The earliest press release I saw for Nanochip is dated 13 January 1998. It says: "Within a year, the company plans to introduce a disk-drive replacement for portable applications in the form of a chip-sized component that holds 250 Mbytes of data." This stuff is harder to do than even the insiders think it is to do.

The media for the Nanochip is "unique charge storage material." A 12-volt pulse stores charge to write a bit and a negative 12-volt pulse erases a bit. The read/write heads operate like small scanning-tunneling microscopes. For reading, a small voltage is applied to the read/write head (the media is grounded). The distance between the read/write head and the surface is adjusted as necessary to maintain constant current. This height above the surface indicates whether the read/write head is over a one or a zero.

The stationary array has 400 read/write heads. The read/write heads are organized ten to a "platform" with forty platforms in eight columns and five rows. Maximum access delay is 0.5 ms. The data transfer rate is 24.8 Mb/s.

Electrostatic actuators move the media in the horizontal plane.

Other projects. HP has experimented with atomic resolution storage using phase-change materials capable of storing 1,000 Gb/in^2. Hitachi has experimented with field evaporation, which selectively removes atoms from the media to represent bits. Hitachi's method could store millions of gigabits per square inch. Canon has experimented with conductance change to store bits. Kionix, Cornell University, U.C. Davis, Georgia Tech, and others have experimented with MEMS-based storage.

State of MEMS-based storage

It's obvious from what you have read that no one is in volume production of MEMS-based storage. So why bother explaining it to you? I'm explaining it for two reasons. First, the growing market for portable devices will demand a MEMS-based storage solution. Second, lessons of MEMS-based storage can be applied to other MEMS applications.

MEMS-based storage is immature. Everyone seems to agree that there will be moving media on one chip and an array of read/write heads on another chip. Beyond that there is lots of experimentation and there are no clear choices for media, for motion actuators (principally electromagnetic or electrostatic), for read/write heads, and for array scanning and data organization. Essentially every important decision is still to be made.

The state of MEMS-based storage is much like that of IC macros—logic building blocks—just after the invention of the integrated circuit. In the early days of the integrated circuit, manufacturers introduced families of IC macros identified by acronyms describing their internal construction (TTL, DCTL, DTL, RTL, ECL, I^2L, etc.). Manufacturers knew that IC macros were a good idea, but were fishing around for the "right" choices to dominate the market. TTL (transistor-transistor logic) emerged as "the solution" for IC macros. It's too early to know the winners in MEMS-based storage, but we can make educated guesses.

Media and read/write heads. The media and the read/write heads are closely related. Media shows the widest range of experimentation, from ordinary magnetic coatings to manipulating individual atoms. The likely near-term (five years) winner for MEMS-based storage is magnetic media. Magnetic coatings are familiar and they benefit from forty-five years of development. Design of the read/write heads for magnetic media benefits from the same familiarity and development history. Read/write heads for magnetic media have an important benefit: they do not have to contact the media. Media and read/write methods that require the heads to be in contact with the media surface will probably prove insufficiently reliable. The long-term winner (perhaps ten years out) may be atom-scale (though I can't guess which variety).

Actuators. Actuators I looked at were all either electrostatic or electromagnetic, though others (linear-drive motors, for example) are possible. Electrostatic actuators are likely to dominate because they are simpler and are more compact.

Array organization/data management. I didn't catalog storage array organization and data management because it would be too much detail. It is, nevertheless, an important and unresolved issue. Hard disks have a tracks and sectors model that doesn't have an obvious analog when the media moves in X and Y directions under an array of read/write heads rather than simply rotating under a single read/write head.

History of magnetic storage limits

A 1970 paper published in the *Journal of Applied Physics* posited an ultimate density prediction for hard disks. By 1976, IBM's 3350 "Madrid" hard disk had exceeded the limit's density. "Mechanical Limitations in Magnetic Recording," published in 1974 in *IEEE Transactions on Magnetics*, set a new limit. The next IBM hard disk, the 3370, exceeded that limit in 1979. I've probably lost a few limits on the way, but since the 1974 "ultimate density" limit fell, the superparamagnetic limit has been set at 20, 40, 50, 100, and now 100-200 Gb/in^2. So, don't put money on the proposition that 100-200 Gb/in^2 *is* the real limit.

Nanomagnetics. Nanomagnetics (www.nanomagnetics.com), a startup working within England's Bristol University, plans to push back the superparamagnetic limit with a unique approach. Researchers at Nanomagnetics are exploiting self-organizing properties of the ferritin protein.

Iron is the central atom of the heme group, a protein group (e.g., hemoglobin, myoglobin, cytochromes) that captures oxygen in the lungs and transports it to the body's cells. Ferritin is your body's iron-storage protein. The ferritin protein provides buffer storage and controls the release of iron. Twenty-four peptide subunits self-assemble to make the ferritin protein—a hollow sphere with channels (for the iron to enter and exit the hollow center of the molecule). The ferritin protein forms a 12-nm sphere with a central cavity of 7.5-8 nm (see fig. 4). Iron, in the form of the mineral ferrihydrite, attaches to the inner wall of the sphere.

Nanomagnetics removes the iron in the protein's center to form "apoferritin." It then grows grains of a platinum alloy (e.g., CoPt or FePt), which has excellent recording properties, inside the apoferritin's core. The apoferritin's cavity confines and controls the size of the platinum alloy particles formed. Nanomagnetics is working on coating the modified ferritin proteins with inert silica jell that can be spin-coated onto a substrate. The ferritin proteins, with their magnetic domains form a hexagonal close pack (HCP) on the substrate (see fig. 5).

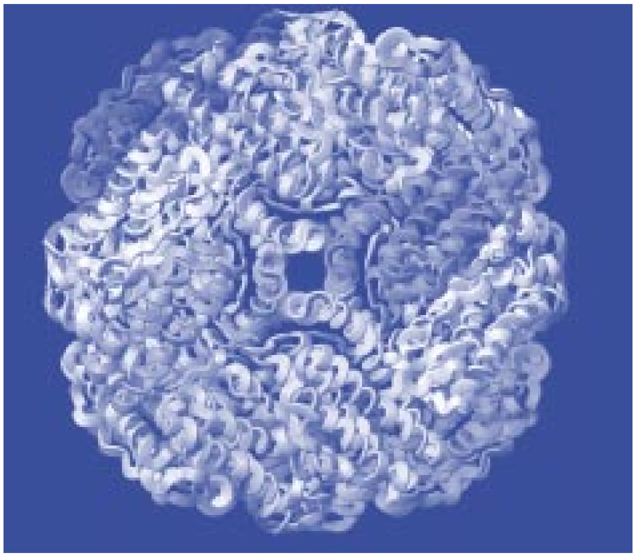

Fig. 4. The ferritin protein self-assembles into a hollow sphere from twenty-four identical subunits.

Each ferritin molecule encapsulates and isolates a magnetic particle. Magnetic domains are the dark spots in fig. 5. The light-colored ring of a ferritin molecule surrounds each magnetic domain. These magnetic particles can be aligned either for longitudinal or for perpendicular recording. The theoretical bit density for this arrangement would be an astounding 4,500 Gb/in^2. (Today's superparamagnetic limit is 100-200 Gb/in^2.) The diameter of the ferritin protein's 12-nm sphere is about a *tenth* of leading edge semiconductor-process line widths, which are about 130 nm. Engineers are beginning to harness biological structures in the service of electronics applications.

Producing magnetic coatings with terabit resolutions is not sufficient to forecast huge near-term gains in hard disk capacity—there remain enormous challenges in designing read/write heads and control mechanisms able to exploit these media improvements. These tasks are more suited to the micron-scale world of MEMS-based storage than they are to the hard disk's macro world.

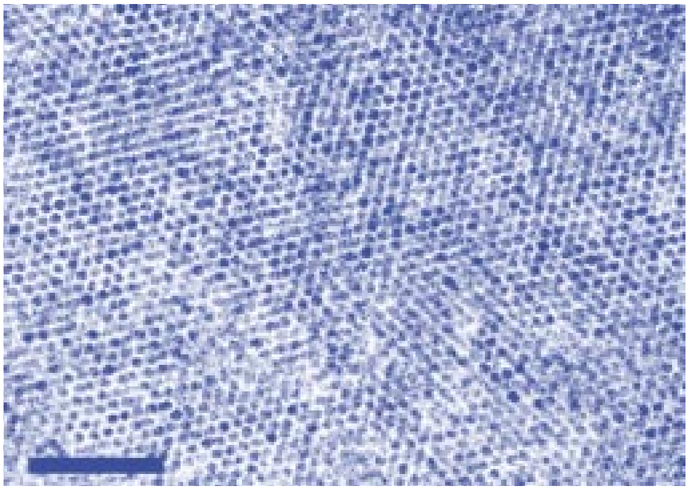

Fig. 5. The ferritin composite film: light circles are the ferritin protein; dark spots are the magnetic domains. The scale bar is 100 nm.

Lessons from MEMS-based storage

The Cahners In-Stat Group projects that the market for MEMS in consumer electronics will climb from $200 million in 2000 to more than $1.5 billion in 2005. Are MEMS ready for high-volume production? For MEMS-based storage, the answer is "no." For MEMS-based accelerometers, the answer is "yes." **Analog Devices** (ADI) produces 700,000 MEMS accelerometers per week at about $4 each. So, how do we know which applications are ready for high-volume production and which are not?

Today's rotating-media hard disks serve today's high-volume markets well, but are not ideal for the emerging market in portable devices. MEMS-based storage is a good match for portable devices, so there is strong incentive to develop MEMS-based storage. MEMS-based storage will eventually have compelling advantages in size, in durability, in power savings, in storage density, and in performance. Many companies, seeing these potential advantages, are developing MEMS-based storage solutions. Today, however, there are still too many unsolved problems.

Industry observers keep thinking that the superparamagnetic limit or some other barrier will stop or slow improvements in hard disks, but the hard disk industry has proven itself amazingly persistent. Instead of slowing, hard disk improvements have accelerated in recent years, pushing back the superparamagnetic limit. Nanomagnetics shows that there are options for pushing the superparamagnetic limit back by orders of magnitude as engineers adapt biological solutions to the world of electronics. MEMS-based storage systems, which are inherently micron scale, are more suited to manipulating terabyte per square inch storage arrays than are macro-scale hard disks. That makes MEMS-based storage the long-term winner, but as MEMS invades the conventional hard disk, hard disks will continue to scale cost-effectively. Samsung is probably close with its estimate of $3 billion by 2005 for MEMS-based storage. That makes MEMS-based storage a small market relative to a hard disk market more than ten times its size.

How can MEMS take off if the whole semiconductor industry is sinking?

Whining about the semiconductor industry is everywhere. I've even done my share. Tech stocks are down. Layoffs abound. Everyone seems to be missing last quarter's forecast and is cutting estimates for the future. Dataquest recently projected worldwide chip sales would reach only $188 billion—a shrink of 17 percent from last year's $226 billion. Is the whole semiconductor industry circling the drain?

The spiking line in fig. 6 plots annual growth rate for the semiconductor industry between 1981 and 2005. Historical data (and part of the forecast) comes from the Semiconductor Industry Association's web pages (www.semichips.org), but I've updated the estimates for 2001 to 2003 with information from Dataquest.

The spiking line in fig. 6 is the annual growth rate of the semiconductor industry. No two years are the same. The industry has grown as much as 47 percent in a single year (1984), and it has shrunk as much as 17 percent (1985). The oscillating line in fig. 6 is the moving five-year growth rate. It varies, but it's always positive. The relatively smooth line in fig. 6 is the cumulative growth rate.

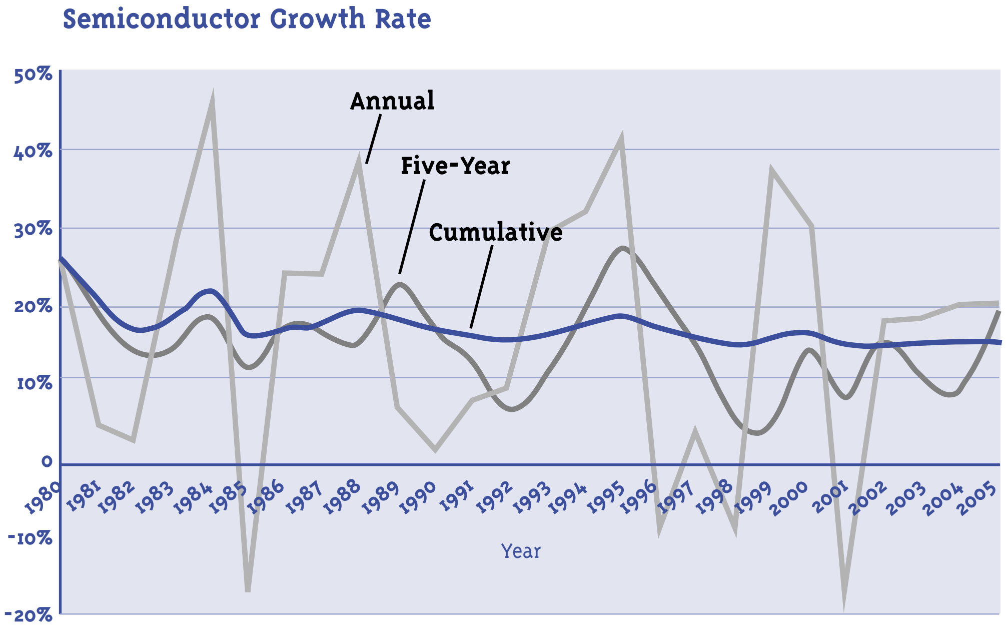

Fig. 6. Semiconductor growth rate by year (spiking line). Forecast decline for 2001 is similar to the decline in 1985.

The cumulative growth rate is about 16 percent and has never fallen below 15 percent. This year's projected decline in the growth rate is about the same as 1985. After fifty years of double-digit growth, perhaps the semiconductor market has run its course and has saturated. Is it time to look elsewhere for investment?

The decline in 1985 followed a year of *47 percent* growth. This year's projected decline follows two consecutive years of more than 30 percent growth. In years of high growth, manufacturers grow as fast as they are able in order to meet demand. Grow with demand or lose market share. In times of high growth, OEMs (original equipment manufacturers, companies that buy components and build systems) accumulate component inventory as they ramp production. OEMs overbook with their suppliers, placing double and triple orders, to ensure that they will get the components to build systems for the expanding market. Component manufacturers ramp production to meet this (inflated) backlog. When the market slows, OEMs are stuck with huge inventories and component manufacturers see orders evaporate. Ron Wilson's 8 January 2001 *EE Times* editorial labels it the "Cycle that won't die."

As long-term investors we need not be concerned with the "cycle that won't die," but how do we distinguish between a down cycle and market saturation? After all, last year's $226-billion semiconductor market was more than ten times its 1985 value. In each of the last two years, it grew by more than twice its *total* value in 1985. Has the semiconductor business reached maturity? Not by a long shot. The Semiconductor Industry Association's 2000 update of the *International Technology Roadmap for Semiconductors* (http://public.itrs.net) forecasts Moore's law progress through 2014. Moore's law semiconductor improvements expand the range of applications. Semiconductors have barely begun to invade some sectors. For a look at how semiconductors will invade the automobile, see "The Silicon Car," in our sister publication, *The Huber Mills Digital Power Report* (December 2000). Digital cameras and high-resolution inkjet printers will transform chemical processing of film to pollution-free digital processing. Semiconductors are just beginning their massive invasion of toys.

The semiconductor industry will continue to thrive and it will continue to grow. Chips shrink to use less power, enabling additional portable applications. Higher-performance new designs capture applications that were once out of reach. Old designs get cheaper, invading applications for which they were once too expensive. New designs offer compelling advantages in capability, efficiency, cost, and weight.

Nick Tredennick and Brion Shimamoto
17 May, 2001

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 and Xilinx (ALTR http://www.altera.com) (XLNX http://www.xilinx.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 and Xilinx are positioned to be major suppliers in tethered applications such as the base stations that support 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.

Calient (* http://www.calient.net)
Calient is a pre-IPO startup that builds photonic switches for the all-optical network core. It builds its own MEMS components. Calient has expertise in MEMS components in Ithica, NY through its acquisition of Kionix and through its own experts in Santa Barbara and San Jose, CA.

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.

Triscend (* http://www.triscend.com)
Triscend builds microcontrollers with configurable peripheral functions and with configurable inputs and outputs. Triscend helps consolidate the microcontroller market into high-volume, standard chips.

Technology Leadership	Company (Symbol)	Reference Date	Reference Price	4/30/01 Price	52-Week Range	Market Cap.
General Programmable Logic Devices (PLDs)	Altera (ALTR)	12/29/00	26.31	25.29	18.81 - 67.12	10.3B
Dynamic Logic for Mobile Devices	QuickSilver Technology, Inc. (none*)	12/29/00				
MEMS Foundry, Dynamic Logic	Cypress (CY)	12/29/00	19.69	22.60	13.72 - 56.63	2.9B
RF Analog Devices, MEMS, DSPs	Analog Devices (ADI)	12/29/00	51.19	47.31	30.50 - 103.00	16.6B
Configurable Microprocessors	ARC Cores (ARK**)	12/29/00	£3.34	£1.18	£0.75 - 4.65	£499M
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	2/28/01	38.88	47.47	29.80 - 98.31	14.3B
Configurable Microcontrollers (Peripherals)	Triscend (none*)	2/28/01				
Silicon for Wireless RF, GPS	SiRF (none*)	12/29/00				
Microprocessor Instruction Sets	Transmeta (TMTA)	12/29/00	23.50	17.44	10.67 - 50.88	1.7B
Photonic Switches	Calient (none*)	3/31/01				

*QuickSilver, SiRF, Triscend, and Calient 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.