Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

The Next Digital Revolution: A New World View

A look at the real world gives us an idea of changes on the way in electronics. Since it's convenient, let's start with my office. It's on the second floor of a small building near my house. The window in front of my desk looks out on redwoods, oaks, and madrones. The house is to the left, across the driveway. The garage is down the hill a hundred or so feet away. I see leaves and limbs, depth and color, sky and earth. The amount of information available in this scene is so large that it is difficult to estimate. Let's take a shot at it by starting with the computer screen in front of me. An array of twenty-five of these two-megapixel displays would effectively block my view. Is the scene outside my window, therefore, fifty megapixels? No.

At seventy-two dots per inch, the puny two-megapixel display has no true depth or fidelity. The unaided human eye can resolve about 1,000 to 3,000 dots per inch, which means I can distinguish features about a thousandth of an inch wide. Of course with the aid of a microscope I can zoom in on a tiny portion of a single leaf and discover that it has features much smaller than a thousandth of an inch. So it bears much more than a thousand bits of information per inch. In fact, it bears so much information that with the right government grant I could spend the rest of my life examining the features of a single leaf. And then I could haul out the electron microscope. And then I could… Ok, I won't do that. Let's just say, for simplicity's sake that the scene has fidelity to ten nanometers, about a hundredth the width of a human hair, and a tenth the width of the circuit lines in a leading-edge microprocessor in 2002. That's about 2.5 million dots per inch, or almost 70,000 megapixels for the scene, about 35,000 times as much as my computer screen can display.

My computer, and the network to which it is linked, works very hard to get me the information on my screen. Every year computers process more and more bits, faster and faster, using more and more power in an effort to bring me crude approximations of the real world, like the scene outside my window. Wow, after fifty years the computer can deliver a wedge of information 1/35,000 as good as I can get outside my window—if the network can deliver it! Meanwhile, how much effort is a tree in this scene exerting to get me all that information? None. The tree's not doing anything to specifically communicate with me. The tree doesn't have to be super smart; it doesn't worry about protocols, bits per second, storage capacity, or power. And it sure isn't very fast. It's just standing there minding its own business. The sun is illuminating the scene and my eyes are the beneficiaries. I sort the available information for the parts that interest me.

My computer is dumber than a tree. OK, that's not really fair. My tree cannot add, subtract, divide, multiply, or plot a course to the moon. And as for fast Fourier transforms, you can just start by forgetting "fast."

My computer is “dumb” in the sense that it is not very good at comprehending, or interacting with, the real world, the rich analog world of color, sight, sense, and sound. As powerful as the computer is, until it can get a lot better at interacting with the world outside my window, or even being able to function in that world, it will never fulfill the promises that have been made on its behalf over the last fifty years.

If the computer ever masters the ability to interact with the real world, directly and without extensive human mediation, then “the computer age” will be even more transformative and revolutionary than we have yet imagined. And the microchip industry, enabling this change, will be reborn, achieving growth rates, profile, and a share of the national wealth dwarfing the achievements of the past thirty years.

That is what is happening now. The second microchip revolution is beginning right now with dramatic developments in the microchip industry, which we sum up in the term *dynamic silicon*. The world of dynamic silicon is the world of chips whose *circuits* change logically and that actually *move* physically, to enable the computer to function in and interact with the real world. If, as George Gilder has written, the true meaning of the computer age is the “overthrow of matter” by injecting the power of mind into every human activity, however apparently physical, then dynamic silicon, which joins the computer to the world, is the consummating force in that vision.

To see how dynamic silicon is a consummation of the computer’s evolution, let’s revisit some history, starting with the mainframe.

Queen ant mainframe

The mainframe computer is like the queen in an ant colony. Privileged attendants nourish and maintain it, while hoards of workers stockpile the nourishment. I should know, I was one of the workers. I became an attendant for a while, but now I’m a worker again.

When I was in college, the computer was a big thing—literally. At my school, the computer had its own building. Inside the building, there were areas where we worker ants sat and programmed and there were areas where we punched the cards that fed the computer. The computer lived in a special raised-floor room behind a wall of glass.

When I thought my program was ready, I delivered a box of punch cards to an attendant on the other side of the glass wall. The attendant took the cards to the computer. A half-day or so later, the attendant returned the box of cards together with the computer’s results.

This system worked for the ants for millions of years. The evolution of electronic systems, however, is millions of times the rate of evolution for biological systems. Within a few (computer) generations the minicomputer arrived to change the model. Smaller and cheaper than the mainframe, the minicomputer could be fed directly by worker ants with no need for the glass house or the attendant priesthood.

With the PC, once again, the model for human-computer interaction changed. Each worker gained a personal computer to feed. In a few more generations, the computers became mobile, sort of, and could accompany us, rather than waiting for us to come to them.

The evolution of the computer is driven by the astonishingly rapid development of its brain, the central processing unit or CPU. In the mainframe era, the CPU dominated the computer system, weighing thousands of pounds and occupying most of the computer room’s volume. Today’s CPUs are smaller and lighter than the smallest coin in your pocket, inconsequential to the weight and volume of the personal computer, but millions of times more powerful than the early mainframe CPU.

In fact, progress in computers is so rapid that everything that has gone before, all the wondrous marvels of the computer age, are dwarfed by what we can do today. Computer power (as measured in transistors on an integrated circuit) still doubles every eighteen months. That’s a compound annual growth rate of sixty percent–and it’s been growing at that rate for thirty years! Take a close look at the following chart where I display the growth in the number of transistors on an integrated circuit. The growth is so big its graph is usually done only in log scale but we are showing the more dramatic version to get the point across. The number of transistors on a chip grows by a factor of ten *every* five years.

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 800.229.2573, e-mail us at dynamic@gildertech.com, or visit our website at www.dynamicsilicon.com

The Future of Microprocessor Performance

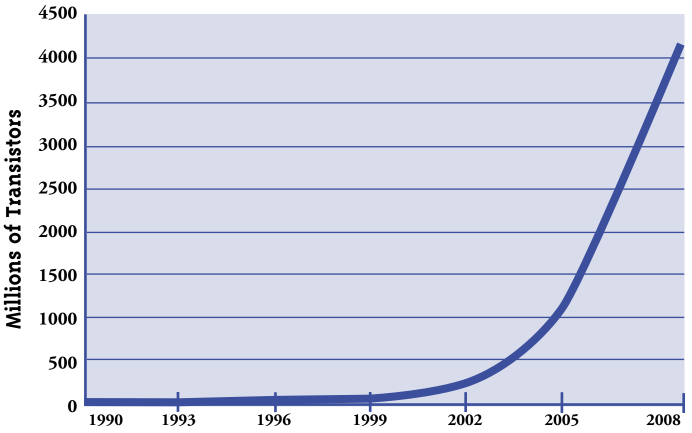

This chart shows progress in the number of transistors on a microprocessor (about a million in 1990). It's plotted on a linear scale to emphasize a point: the on-chip capability we will have in the near future dwarfs the entire history of developed capability.

This graphs the change in your computer's capabilities. But it is not a graph of the progress we have made in how we use the processing capability.

The problem is, in contrast to the phenomenal rate of evolution in computer systems, we humans don't change at all, which may be why we're more comfortable battling ants than computers. If humans and computers are elements in a social system, we're standing still (from the computer's point of view) and the computer is rapidly evolving its raw capabilities. We set the rules for human-computer interaction in the late 1940s, and in the beginning the computer didn't do anything for itself. We wrote the programs, we put the programs and the data in the right form, we delivered the programs and the data to the computer, we collected the results, and we interpreted the results.

Even as the computer has shrunk in size and grown in power a million fold, we have retained our original model of interaction with it. We're still dealing with programs, programming languages, protocols, bits per second, frequency, power, instructions per clock tick, and bus width.

That's about to change. Soon the computer will be capable of collecting and interpreting its own raw data and it will be capable of more direct action than to print results for the human to interpret. *Originally, humans took the problem to the computer. Today the computer is small enough to be taken to the problem.* That's a big change, but it isn't nearly enough. We still treat the computer as if it belongs in a room with a raised floor. We connect to it with protocols and programs and graphical user interfaces—but these should be artifacts of a bygone time. Instead of a few queen ants, we have many. The computer processes everything in discrete (digitized) chunks and it works on these chunks one at a time. The world, meanwhile, is a continuous realm where things happen all at once. That's why computers need to run at a zillion megahertz just to handle the data from my lazy old tree. At any instant, they are looking at only a tiny piece of the problem.

The computer's world divides

The computer's world of bits and bytes, programs and protocols isn't going to die; it will continue to grow. Computers will divide into those that work with bits and bytes and those that work with the real world. The computers that work with bits and bytes will connect directly to the information and power grid—the Fibersphere and the Powercosm—with fiber and wires. This vast, distributed network, dubbed the "global information grid" by the U.S. Department of Defense, comprises the world's combined resources of computing, access ports, data transport, and storage. Thanks to the Internet, we can use the world's accumulation of information and computational resources from any convenient access port. It's magnificent, but as a computing paradigm, it's more of the same, an extension of where we've been headed for thirty years or so. And we know how to do it, so breakthrough market-beating investment opportunities are increasingly rare.

Lend me your ears, your nose...

If we're to get away from the concepts of bits per second, bus width and protocols, and of having humans as an integral part of the data collection process, then mobile devices will have to become more active and more direct in collecting real-world data. To become better at interacting with the real world and with humans, mobile devices will need their own eyes, ears, noses, and other senses. Some of these devices will have to see and hear. Some will need to know where they are and what's happening around them. Some will have to sense motion.

Even today, without dynamic silicon devices, we could build a computer with some ability to deal directly with the world. We could add optical sensors, audio sensors, pressure sensors, motion sensors, chemical sensors (electronic tongues), and gas sensors (electronic noses). Unfortunately, by the time we added all that stuff to interact with the world we would have to leave it at home, since it would be too unwieldy and require too much muscle power to go most places.

Dynamic silicon will solve that problem. Once again, the semiconductor will save the day. The same tiny integrated circuits that use computation and logic to digitally recreate analog functions can perform those functions directly. Integrated circuits can sense pressure and temperature and can even have *moving parts*. Called microelectromechanical systems, or MEMS, these integrated circuits that interact with the real world have been under development since the mid-1950s, but progress has been slow, particularly for MEMS with moving parts. That is changing now. We are seeing breakthroughs that will transform the industry and our world in the next few years.

MEMS are what they sound like: tiny electromechanical systems: turbines, motors, gears, moving mirrors, filters, and so on. Semiconductor processes build these systems, but they do not require the expensive fine-line process equipment required for commercially-competitive semiconductors. MEMS do not need a state-of-the-art semiconductor process; a MEMS fab can be a few generations behind. MEMS can leverage the processes and the accumulated expertise that has been paid for by microprocessor and memory applications. Therefore, I believe MEMS will achieve a growth rate that exceeds the historic growth rate of the semiconductor industry.

Today, there are hundreds of millions of MEMS sensors in the field and tens of millions of these MEMS already have moving parts. The airbag sensor in a modern automobile, the first large-scale use of a MEMS sensor with moving parts, uses an integrated circuit with a tiny on-board accelerometer. The automobile may also use integrated MEMS sensors to measure mass airflow, manifold pressure, and fuel pressure, and to analyze exhaust gas. Inkjet printers use tiny passive MEMS (by passive, I mean they don't have moving parts) nozzles to squirt precisely measured droplets of colored ink at the paper. The read head in a typical hard disk is a passive MEMS sensor. In current applications, MEMS with moving parts sample the environment or accomplish mechanical tasks, and they are mass produced with precise tolerances.

More to come

The computer industry with its printers and hard disks, the biomedical/chemical industry with its disposable pressure sensors, and the automotive industry with its airbag accelerometers and engine sensors are today's volume applications for MEMS. Soon other applications in automotive/industrial (including automotive, business, industry, and computer), optical, biomedical/chemical, and radio frequency (RF) electronics will follow as voracious consumers of MEMS. As the market for MEMS develops, applications will move from the wired world to the mobile world. MEMS applications in the wired world (automobiles, backbone networks, medical instrumentation, et al.) will pave the way for mobile applications. Crude pressure, chemical, and gas sensors in wired applications are the ancestors of MEMS-based sensors in mobile devices. The next breakthroughs will come in areas including optics, biomedical and chemical analysis, and RF electronics.

Optical MEMS

MEMS are the basis for one of the most powerful breakthroughs in fiber-optic networks. Conventional switches in an optical network convert the light to electrons for routing and convert back to light for transmission. The electronics in the middle of the switch limits the transmission to preset bit rates and protocols. By contrast, the MEMS-based, all-optical switch, essentially a tiny moving silicon mirror, directs its light across free space to its destination fiber.

Coming from companies such as **Calient**, **C Speed**, **Cronos** (**JDS Uniphase**), **Xros** (**Nortel**), **Lucent**, and **MEMX**, micromirror-based, all-optical switches are already capable of switching multiple terabits of information across hundreds of ports—thousands soon—regardless of bit rate or protocol. Similarly, MEMS will provide configurable add/drop multiplexers to add or drop signals to generate thousands of different frequencies of light. Multiplexers and de-multiplexers employing MEMS-based tunable lasers will direct multiple light streams onto one fiber. The network's amplifiers will employ MEMS-based attenuation, dynamic gain equalization, and dispersion compensation. The phone companies spent decades eliminating moving parts and electromechanical switches from their networks. But now moving parts are coming back, at the same micro scale and with the same reliability as digital chips, but recapturing the simple physical efficiency of mechanical systems.

Biomedical/Chemical MEMS

Biomedical/chemical MEMS will create enormous wealth for companies able to navigate the hazardous maze of biology, chemistry, and electrical and mechanical engineering, to satisfy the demands of the biomedical/chemical practitioners, and to satisfy the federal regulatory bureaucracy. To give you a feel for the size of the pot of gold at the end of this rainbow, I'll cite two examples.

Batch-fabricated, like microprocessors and memory chips, MEMS can be cheap to the point of being disposable. There are over 150 million diabetics in the world today, 15 million in the U.S. alone. In treatment, diabetics typically sample their own blood to test its insulin level, and adjust as necessary. Recommended sampling rate is four times a day, but the process causes pain, discouraging compliance.

Kumetrix (http://www.kumetrix.com), a pre-IPO startup in Silicon Valley, replaces the typical lancet and blood-testing strip with a battery-powered, handheld diagnostic instrument using disposable MEMS cartridges, consisting of a microneedle with a cross section smaller than a human hair, a silicon cuvette (a chamber for the blood sample), and a small window. This device uses capillary action to painlessly draw approximately 100 nanoliters of blood into the cuvette where it is mixed with a chemical reagent. The diagnostic instrument's laser then measures and displays the blood's glucose concentration. In semiconductor manufacturing, a single six-inch diameter silicon wafer (at least a generation behind the state of the art for microprocessors) yields several thousand sampling devices. With tens of millions of diabetics testing their glucose levels several times a day, the potential worldwide market for this MEMS application is tens of billions of units a year. Most diabetics would prefer a sterile mosquito bite to a stab wound.

Today's DNA testing requires a room full of sophisticated chemical analysis instruments. Soon, MEMS will reduce this to a portable instrument incorporating an on-chip chemical analysis chamber. Polymerase chain reaction (PCR), genetic analysis, begins with a DNA fragment taken from a single cell. This fragment is too small for direct analysis, so it is placed into a solution and thermally cycled. When the solution is heated to 92°C the DNA fragment's double helix splits into two complementary halves. When the solution is cooled to 65°C the individual strands construct duplicate fragments, doubling the DNA's concentration in the sample. Twenty to thirty thermal cycles bring the concentration in the sample to a level suitable for reliable testing. Since the amount of solution required for conventional analytical instruments is large, the solution has a large thermal mass. This means the heating and cooling cycles required to increase DNA concentration in the sample can take a long time.

A MEMS thermal-cycling chamber reduces the thermal mass to about fifty microliters, reducing thermal cycling to a tenth of the time. In addition, the small amount of solution reaches the required concentration in fewer cycles and substantially reduces reagent use. **Cepheid** (CPHD, http://www.cepheid.com), another Silicon Valley company, builds MEMS-based instruments for rapid thermal cycling and analysis of small nucleic-acid samples.

MEMS in RF electronics

Radio-frequency signals enter and leave a cell phone through the antenna. The back-end of the cell phone, which houses all the digital processing functions, is state of the art (more on this later). It has benefited from thirty years of integrated-circuit development. Between this digital back-end and the antenna is the RF electronics. The RF electronics converts analog signals on the antenna's input into bits and bytes for the cell phone's digital back-end CPUs and ASICs (application-specific integrated circuits) to process. The RF electronics also converts the back-end's output bits and bytes into the RF signal to transmit. If you open the phone and look at the RF electronics, it's a mess. It's a tangle of discrete components wired and soldered together as if it were still the 1950s and nobody had ever heard of the integrated circuit. Compared to the electronics in the back-end of the cell phone, this is Stone Age technology.

> **Once we've mastered obvious "large physics" applications, we'll begin to understand and to exploit MEMS more in their native "small physics" environment.**

Inductors, transformers, relays, and moving-plate capacitors cannot be fabricated on a planar silicon integrated circuit, i.e., a microchip. Therefore, these basic elements remain implemented as discrete components and not as integrated circuits. The microchip has so many advantages in size, cost, reliability, ease of manufacturing, and eliminating the assembly of discrete components, that electronics engineers avoided discrete components, instead aping their functions with elaborate substitute circuitry that could be fabricated on a chip. Performance suffered, but if performance parameters were not critical the savings justified the sacrifice. And when, say, an inductor or moving-plate capacitor was absolutely required—as in the RF world—an old-style discrete component was attached externally, which explains the mess in your cell phone.

MEMS are microchips. And all those discrete components we surrendered on the microchip—or soldered externally—can be made as MEMS. With MEMS, which are not necessarily planar and which can have moving parts, we get back much of the quality we sacrificed for the cost savings and convenience of microchips. The Q-factor (a measure of quality) of a discrete-component coil (a typical Q value might be 50 to 150) is much better than the circuit we use to imitate the same function on a planar integrated circuit. A MEMS coil recovers the quality lost in giving up the discrete coil. A true on-off switch is impossible outside the mechanical realm. That may come as a shock since we all know that transistors are switches. But transistors leak. They are never entirely on or entirely off. The leakage is so small that it does not affect their function in a computing device. Leakage in the off state means, for example, that they cannot isolate a high-gain amplifier:

An introduction to nanotechnology

MEMS and nanotechnology are often lumped together by the popular press. I suppose this is because they are small, equally magical to someone not versed in technology, and have moving parts. MEMS and nanotechnology are vastly different. The differentiator is scale. Today's integrated circuits are designed with geometries approaching 100 nanometers (a millionth of a millimeter). At that scale, lines are about a thousand atoms across. MEMS devices use geometries of about 1000 nanometers, so lines are about ten thousand atoms across. Nanotechnology uses geometries on the scale of the atoms themselves (about ten atoms fit in a nanometer). We know how to build semiconductor structures at 100 nanometers and we're way into knowing how to build electromechanical structures at 1000 nanometers (a micron), but we don't have experience or much knowledge about building atom-scale structures.

MEMS and nanotechnology are separate areas. MEMS designers build (trailing-edge) semiconductor-scale machines. Nanotechnology is the theory of building atom-scale structures. The ideas and range of applications may be similar: small machines, but the means to achieve the result are substantially different. MEMS exploits the use of semiconductor processing equipment to make small machines. Nanotechnology must develop its own molecular-manipulation equipment to build its systems. It is not yet a solved problem. Nanotechnology does not ride the coattails of the semiconductor processing industry. Its progress, consequently, will be much slower.

some current will flow, creating noise. And in the on position they waste signal strength.

The transistor switch was invaluable because electronics was the only way to get that small, that cheap. But micro-mechanical switches, micro-sized electrostatic switches with physical gates that actually open and close will replace transistors for some functions, especially mobile devices. The "on" resistance of such a switch (the failure to transmit power) will be negligible compared to a semiconductor switch, and its "off" or open resistance will be infinite. It will be slow compared to a semiconductor switch, but fast enough for RF functions like switching among various antenna elements, filters, and amplifiers to optimize performance.

MEMS summary

I expect high growth for MEMS in automotive/industrial, optical, biomedical/chemical, and RF electronics applications. This isn't the way professional market analysts categorize the MEMS market, but it's the way I think about it from the perspective of mobile devices. We've already got high-volume applications in automobiles (engine sensors and airbag accelerometers), in computers (inkjet printers and hard disks), and in biomedicine (disposable pressure sensors). Production volumes for these applications help the industry go down the learning curve, which further expands the range of affordable applications and new opportunities.

Large physics, small physics

Today MEMS exploit the Lilliputian world to make things we already know how to build such as accelerometers, pressure sensors, magnetic read heads, inkjet nozzles, and electromagnetic relays. I call this the "large physics" world, because these devices are built to function in the world familiar to us all—the one dominated by gravity. But there is another "small physics" world, the world of bacteria, viruses, small insects, and now MEMS where gravity is unimportant compared to forces like surface tension, friction, and electric charge. "Small physics" can explain why the fly seems to have no respect for which way is "up" when it lands on walls and ceilings. Once we've mastered obvious "large physics" applications, we'll begin to understand and to exploit MEMS more in their native "small physics" environment.

DSPs and microprocessors fall short

MEMS in future mobile devices will enable direct interactions with the real world. Mobile devices will sense and manipulate the analog, high-fidelity world. The cell phone is the prototypical mobile device of today. I said earlier that the back-end of the cell phone contained state of the art electronics. It does, but it's the wrong stuff. The core of the back-end of today's mobile devices, such as cell phones, PDAs (personal digital assistants), pagers, and GPS (global positioning system) receivers, is the microprocessor, application-specific integrated circuit, and the digital signal processor (DSP). The microprocessor and the DSP are built for the world of bits and bytes, programs and protocols, and humans; they aren't efficient enough or fast enough for the analog, high-fidelity world. If the sensors and actuators become direct and more efficient, the back-end's logic has to become more direct and more efficient. If a back-end based on the DSP and microprocessor won't do, then there'll have to be a change; and where there's change, there's opportunity.

The microprocessors and DSPs in today's mobile devices have two fundamental problems: efficiency and performance.

The microprocessor's inefficiency derives from the culture that creates it. Gods of computer architecture somewhere at **IBM**, **Intel**, **Motorola**, **ARM**, **Sun**, or **MIPS** design the instruction set and programming model for each microprocessor. Gods of computer science, with no working connection to the gods of computer architecture, design programming languages. Gods of software, also working independently, write compilers and operating systems for the programming languages and for the microprocessors, respectively. We bring the microprocessor, the programming language, the compiler, and the operating system together to build systems. The system's inefficiency shouldn't be a surprise.

In solving a real-world problem I select a microprocessor, a programming language, a compiler, and an operating system. I must consider the microprocessor's instruction set in selecting an algorithm for the application. It wouldn't do to select an algorithm that required floating-point operations if the microprocessor didn't support them. I then map the algorithm into the programming language, which may not possess perfect constructs for all the requirements of the algorithm. I run the program through the compiler, which may not provide an ideal mapping of language constructs to microprocessor instructions, to produce object code for the microprocessor. A less than ideal operating system provides services to the application program as it runs.

It is a Rube Goldberg contraption that loses efficiency at every step. For the wired world it all makes sense. The scarce resource is the engineer's design time. We're willing to lose efficiency in the design if the designer can be more productive. Designers are more productive programming in C than in assembly language, but programs are more efficient in assembly language than they are in C. In the wired world, we trade efficiency in the end result to achieve higher levels of abstraction, which conserve the designer's time. It's the right choice for the wired world, but it won't do for mobile devices.

If efficiency is poor, performance may be disappointing. Historically, the microprocessor has improved with improvements in semiconductor fabrication. The primary contributor to the microprocessor's improved performance has been higher clock speed. If the microprocessor needed more performance, we cranked up the clock. Double the clock speed, and roughly speaking you double the performance. There is, however, a side effect. Power dissipation also increases directly with clock frequency. Double the clock speed, you double the power dissipation. Up to a point, this isn't a problem for a system getting its power from a wall socket. But it presents the microprocessor and the DSP with a dilemma for future mobile devices. Consumers will demand both high performance and long battery life from these devices.

Before the microprocessor, engineers mapped their applications directly into the hardware (no compilers, no operating systems, just write the equations for the transformations, design the state sequencer, and build the hardware). Input signals were converted directly into desired outputs by hardware built to implement the appropriate transformation equations. Mobile devices of the future will demand the efficiency of direct implementations, but will still demand the flexibility now provided by software.

Direct hardware implementations are like ASICs; they aren't flexible. What is needed is what I call dynamic logic. Dynamic logic is a distant relative of programmable logic devices (PLDs) or field-programmable gate arrays (FPGAs) manufactured by companies such as **Altera** and **Xilinx**. Altera and Xilinx build programmable logic devices for general-purpose applications in the wired world. Dynamic logic is programmable logic custom-tailored to suit a narrow range of mobile applications. Dynamic logic "pages" hardware functions into the system as they are needed.

Microelectromechanical systems will change the world in ways that are beyond forecasting today, because the field is beginning its time of rapid growth.

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.

QuickSilver Technology, a pre-IPO Silicon Valley start-up, 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. QuickSilver calls it an Adaptive Computing Machine (ACM).

Logic for each of the cell phone's protocols and functions can be "paged" into the chip's programmable logic, eliminating the need for a digital signal processor, for ASICs, and possibly even for the microprocessor. Functions that are not paged into the chip's gates do not use power.

Efficiency improves because the implementation is more direct for each function than it is in a DSP-based implementation. The DSP-based implementation runs programs on a fixed set of resources (arithmetic units, shifters, multipliers), giving up efficiency for the sake of simplifying the programming and the hardware resources. The dynamic-logic solution gives up efficiency in "paging" functions into the programmable logic. QuickSilver's bet is that paging the circuit functions onto the chip will use less power than having circuitry that is always resident, but mostly idle.

QuickSilver's Adaptive Computing Machine can be significantly more efficient than a DSP-based implementation of the same functions. Resources (general-purpose logic elements) on the chip can be allocated to the limit of availability for parallel calculation, since the resources are not dedicated to particular functions, as they would be in a microprocessor, DSP, or an application-specific integrated circuit (ASIC). A large fraction of the fixed resources in a microprocessor or DSP may be idle at any particular time. DSPs generally work on data in multiples of a byte. Dynamic-logic implementations can work on any data width (the width can even vary with time to suit the needs of the problem).

As the semiconductor manufacturing process improves, DSPs and microprocessors are built with more fixed resources running at greater clock speeds, so they can tackle ever more complicated functions. But boosting maximum capability does not improve efficiency. By contrast, a dynamic logic device uses resources efficiently by altering its capabilities as needed rather than drawing down the power needed to support latent capabilities. With dynamic logic, QuickSilver may slash power dissipation by fifty to ninety percent compared to a DSP-based system.

New world view reprise

I've covered the divergence of the world into the global information grid and mobile devices. I've given examples of how MEMS and dynamic logic might change the world as they invade mobile devices. Let's get back to the view outside my window. I brought it up to contrast the difference between the open, direct way the real world captures and conveys information and the restricted "communication channels" that the computer uses to capture and to convey information.

As we add an information layer to the real world, we need something better than conventional transmit-receive protocols with communication channels. We need something more like the real world. It's a poor analogy, but the best illustration for this concept is a transponder. A transponder is a transmitter-receiver that, when activated, sends a predetermined signal. A transponder is essentially, an electronic reflector. Today, transponders are used in aircraft and in automated toll booths. Imagine a world where transponders are in our clothing, in our cars, in our carpets, in our porch railings, in our golf balls, and in our shoes. When we "illuminate" an area, it triggers "reflections" from certain places. The "illumination" contains a coded request that selects certain transponders and may even contain the energy to power those transponders.

Where there's change, there's opportunity

Microelectromechanical systems will change the world in ways that are beyond forecasting today, because the field is beginning its time of rapid growth. Among the capabilities I expect MEMS to deliver are the e-nose and the e-tongue. These are cheap, reliable, integrated gas and chemical sensors. In addition, optical MEMS devices, such as integrated mirror-arrays, enable the construction of optical multiplexers that do not require conversion of the light signal to an electrical equivalent. The next generation of MEMS-based inkjet nozzles will eject bubbles smaller than the human eye can resolve. Printers will have the capability to produce prints that are indistinguishable from a chemical-process photo print. This development will spur the conversion of the photo-print industry from its current chemical-processing base to a digital base.

The transition from the microprocessor and from the digital signal processor, to dynamic logic as the anchor chips in the design of mobile devices, will be more wrenching and more difficult than the introduction of MEMS into mobile devices. The microprocessor and DSP are thoroughly entrenched in engineering education, in engineering design experience, in the installed base of development systems, and in huge commercial enterprises based on their continued development and sales. The success of dynamic logic, like the success of the computer, will begin with a compelling application that will provide proof of value to seed the idea in other applications. Watch QuickSilver.

Nick Tredennick Brion N. Shimamoto

Nick Tredennick and Brion Shimamoto
March 2001

Examples of 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.