THE HUBER MILLS digital power report

Powering the 21st Century

July 2002/Vol. 3 Issue 7

Published by Gilder Publishing, LLC

Building the I²C

Next comes the integrated-integrated circuit built with Veeco's tools for the quantum machine shop

Why did it take so long? How did they finally build it? And why have so few people noticed? Last year, Motorola managed to build a gallium arsenide device on a silicon substrate. A new era in the manufacturing of quantum devices will begin in the second half of this year when Motorola and other licensees of its technology start qualifying devices on first generation "evaluation" 12-inch GaAs-on-Silicon wafers.

This is a huge breakthrough, even if almost no one yet grasps how big it is. We've come across only a couple of readily accessible descriptions of the remarkable device that Motorola managed to build—the first went largely unnoticed two years ago in the February 2000 *IEEE Spectrum* article, outlining Motorola's plan. Then another in this past April's *Technology Review*, revealing success. And those two stories hardly mention the extraordinary engineering tools used to build the structure—the atomic-scale molding machines, saws, lathes, sanding wheels, drill presses, and milling machines.

Gallium arsenide (GaAs) is the most widely used semiconductor in the manufacture of extremely fast amplifiers—the ones that are fast enough to boost currents to the gigahertz frequencies used by cell phones and other high-speed radio-frequency (RF) devices. But GaAs is very brittle, and thus difficult and expensive to work with. Ideally, manufacturers would use just enough of it to do the real work, and no more. That would mean building the GaAs circuits on top of a much cheaper and more robust scaffold of silicon (Si). But that's impossible—the lattice structures of GaAs and Si crystals won't fit together, the atomic spacings just don't match. So GaAs circuits are built instead on top of GaAs substrates, which makes them even more expensive.

Motorola got around the problem by depositing a thin layer of strontium titanate ($SrTiO_3$ or STO) on a silicon substrate, and then building GaAs layers on top of it. The future of the cell phone is now a club sandwich: Si-STO-GaAs.

It takes quite a sandwich maker to build it. By comparison with Motorola's triple-decker club, a Pentium is mere white bread and butter. Building an ordinary integrated circuit is quite a feat, but the core tools and techniques are now decades old, and well understood. "Heterojunction" devices like Motorola's use a much more diverse mix of semiconductor materials, and they build much more complex structures. The three layers in the Si-STO-GaAs stack are all atomically perfect crystals, of perfectly controlled dimensions.

There's no reason to stop at gallium arsenide on silicon. If the cell phone can be built on a silicon scaffold, a diode laser and a photo-detector can be built alongside it too, on the same chip, so that a hybrid phone/PDA can communicate optically when there's fiber-optic glass at hand. Logic can go on the same device as well—Pentiums are built on silicon too. And, an array of other sensors might likewise be built on to the same chip, to let the device sense temperature, current, voltage, magnetic field, infrared radiation, ultraviolet light, or perhaps even X-rays. Integrate those additional functions on the same substrate as the digital logic, and you have a *truly* integrated circuit. Call it the "integrated-integrated circuit" or I²C.

To build its Si-STO-GaAs chip, and to secure its nearly 300 patents for GaAs-on-Si technology, Motorola relied principally on Molecular Beam Epitaxy (MBE). The same technology, now configured for mass production of these complex chips, is the key hardware behind "Thoughtbeam," a recently formed joint venture between Motorola and the British provider of semiconductor substrates, IQE. The MBE machines themselves are manufactured by Applied EPI.

Applied EPI was promptly purchased by Veeco (VECO). And with that acquisition, Veeco became the world's largest manufacturer of the most advanced machine tool of quantum engineering, the MBE machine. Veeco started out isolating the rare isotope of uranium essential for the Manhattan Project and was incorporated in 1945. In 1990, a group of the company's executives exercised a management buyout of Veeco's instrument business from the British Unitech, who had acquired it some years earlier. Today, Veeco's $450 million annual sales are anchored in machines that meticulously choreograph the flows of power and material that build films, layers, and junctions of crystalline perfection, atom by atom. The company went public in 1994, and has since made nine acquisitions to consolidate its position as the only vendor of a complete, integrated line of the tools of the quantum machine shop.

Disintegrated Circuits

Invented in 1947, the transistor was the first practical creation of quantum engineering. Physicist and future Nobelist Richard Feynman surveyed the astonishing future that lay ahead in a 1959 speech to the American Physical Society titled "There's Plenty of Room at the Bottom." "What would the properties of materials be if we could really arrange the atoms the way we want them?" Feynman asked. "I can hardly doubt that when we have some *control* of the arrangement of things on a small scale, we will get an enormously greater range of possible properties that substances can have, and of different things we can do."

Though he is often cited as a visionary of (misnamed) "nanotechnology," Feynman was definitely *not* talking about "micro-machines." Micro-electro-mechanical systems (MEMS) are tiny mechanical devices—sensors, valves, gears, mirrors, and actuators—built on to semiconductor chips, along with the electrical circuitry required to interface with them. Build and wire up a suitable set of cantilevers, for example, and you have an accelerometer that can deploy an automobile air bag. MEMS devices reflect marvelous technology and they have a tremendous future in their own right. Some of Veeco's tools are used by the MEMS producers, too. But don't call MEMS devices "quantum," at least not the mechanical parts—they're just plain old mechanical, only smaller. "Micro" though they are, they're still governed by Newton's laws. Quantum devices are very much smaller, and are governed by the fundamentally different laws of quantum physics.

Quantum devices begin with materials that are plucked from favored columns of the Periodic Table. They are then united in an alphabet soup of junctions never imagined or used in any branch of engineering until the advent of the semiconductor diode, and then the transistor. As we have noted before, quantum physics happens in the electron orbits around atomic nuclei—a negative charge (n) situated in quantum proximity to a positive (p) one. A semiconductor junction provides the equivalent—the spherical atom unrolled (so to speak) into a flat n-p layer, that is much larger and that can therefore handle much more total power. The p/n interface *is* the device.

Nothing interesting happens at these interfaces, however, unless they are engineered to atomic-scale levels of precision, just as Feynman described. The only way to make the entire interface perfectly is to build it on fantastically pure and precisely formed crystals, to which may be added minute, precisely calibrated quantities of dopants. Such junctions are extraordinarily difficult to build. But because they operate on atomic-scale junctions, and control power on surfaces, rather than in volumes, quantum devices are blindingly fast, compact, and efficient. Transistors, photodiodes, photo-detectors, and solid-state lasers are quantum switches, emitters, and receivers that dramatically outperform everything that came before.

Further, equally dramatic improvements can be obtained by building many of these devices on to a single crystal, and knitting them together into an "integrated circuit." A Pentium packs 50 million transistors on a single crystal, along with comparable numbers of resistors, capacitors, and conducting connections.

Yet for all that, the integrated circuit as we know it today remains a relatively simple device. It is usually built on an elemental semiconductor—pure silicon—not a compound semiconductor, like gallium arsenide. Gates are built by selectively doping silicon with boron and phosphorus. Oxygen is added to create sil-

Authors	Peter W. Huber Mark P. Mills
Publisher	Lauryn Franzoni
Designer	Julie Ward
Subscriber Services	Laura Parsons
Chairman	George Gilder

The Digital Power Report is published monthly by Gilder Publishing, LLC. Editorial and Business address: 291A Main Street, Great Barrington, MA 01230. Copyright 2002, Gilder Publishing, LLC. Editorial inquiries can be sent to: peterhuber@gilder.com or markmills@gilder.com.

icon oxide insulators. Aluminum (and more recently, copper) serves as a conductor. These same elements are the key constituents of the capacitors and resistors that get "integrated" into the same chip. There are just five key chemical ingredients: silicon, oxygen, boron, phosphorus, and aluminum.

Which leaves a whole lot that isn't integrated into the chip at all. Devices built from that limited list can't handle very high-frequency currents. They can't emit photons efficiently, and they don't make very good detectors in most frequency bands. So radio and optical links are delegated to other semiconductors, on other chips.

In the pursuit of more chemically integrated semiconductors, the first objective is to get more leeway with scaffolds, as Motorola did when it managed to layer GaAs on a silicon substrate. Silicon wafers come in up to 12-inch sizes, twice as large as current GaAs wafers, and working with larger wafers cuts production costs sharply. Silicon is a better conductor of heat, which allows denser packing of higher-power components on the chip. Other substrates—silicon carbide, for example—potentially offer even more attractive properties, in particular for high-power applications (*Quantum Power, May 2001*).

The I^2C chip engineer also wants more leeway with insulators, which are also key components of gates and capacitors. The approach at present is to build insulators by oxidizing the silicon itself. But gates will soon be so thin that silicon oxide will no longer be up to the job—if the layer is built thin enough to maintain transistor performance, electrons will tunnel right through it. It was to address that problem that Motorola initially turned to STO.

The next step is to integrate faster semiconductors capable of handling the very high-frequency currents required for high-speed radios—materials such as GaAs and indium phosphide—the "III-V" compound semiconductors that combine elements from columns III and V of the periodic table. And at least some of these same III-V compounds can also be used to integrate optical functions on to the chip. GaAs, for example, is also a key building block of light-emitting diodes (LEDs) and solid-state lasers. Other candidates for integration would be the compounds that make blue, red, and orange LEDs and lasers: GaN, AlInGaP, and InP.

The I^2C designer will want sensors on his chip, too. Photodetectors—LEDs operating in reverse, to sense light, or perhaps even to receive power (via laser, for example) to power the rest of the chip. Magneto-resistive and piezo-electric materials capable of sensing electrical fields, temperatures, velocity, and even orientation (*A Sense of Power, August 2001*). Integrators may eventually pursue such compounds as mercury-cadmium-telluride (HgCdTe) and indium antimonide (InSb) (for infra-red sensing), amorphous diamond (for ultra-violet detectors), barium-STO (for integrated planer antennas), and perhaps cadmium zinc telluride (CdZnTe) (for X-rays).

For now, the closest anyone gets to this level of integration is in chip-scale packages, which pack multiple die inside a single package (*Packing Power, April 2002*). No one yet builds a single chip that integrates Intel (logic), RF Micro Devices (wireless), JDSU (optical), and Raytheon (IR sensor). But such chips will be built. And when it comes to mass production, there's a good chance they'll be built with equipment supplied by Veeco.

Quantum-Structure Machine Tools

It takes extremely high-grade energy to rearrange matter by the atom. Pound for pound, it takes much more of it to build a Pentium, a laser, or a strand of DNA, than it does to launch a Space Shuttle. The quantum physics at the p-n junction is exquisitely sensitive to the basic chemistry and arrangement of its constituent parts. Atomic-scale engineering requires ultra-pure source materials, exceptionally empty vacuum chambers, exceptionally reliable ovens, ion beams, and molecular beams. Only very precisely controlled power can attain one-in-a-thousand-trillion levels of atomic purity north and south of a junction, while building nanometer-scale structures out of very different elements and compounds.

Quantum engineering machines place atoms on a surface, or take them off, or determine where material will be added or removed. A surface is coated with a protective layer of inert material, which is covered in turn by a mask. The protective layer is then selectively etched off. Then new material is added to the now exposed parts of the original surface. Then the remainder of the protective layer is washed off from the unexposed parts, leaving new vertical structures on the original horizontal plane. The process is repeated again and again, to build up complex high-rises on the original crystalline plane.

Many different physical processes and forms of energy are used in the most difficult phase, the adding of new material to the surface. Doping doesn't add discrete new layers, it injects atoms directly into the existing surface—phosphorus or boron into silicon, for example. Diffusion processes use thermal energy; ion implantation techniques use electrical propulsion. Doping existing layers is hard enough; growing discrete new layers is much harder. In liquid-phase epitaxy, a new layer is crystallized on to the surface out of a liquid solution. Chemical Vapor Deposition (CVD) directs highly reactive gases at the surface, where these precursor chemicals react to deposit a new crystalline layer. Ion Beam Deposition (IBD) accelerates a charged stream of molecules towards the substrate surface. Molecular Beam Epitaxy (MBE) uses evaporation cells, each containing one component material. Each cell has a shutter that opens to allow a specified number of layers of atoms to be deposited, and then closes. Successive layers react and crystallize on the heated substrate.

The importance of MBE can be gauged by reference to what it is used to build

Three other quantum-structure machine tools are still used largely as test instruments or in R&D environments, but are now poised to make the transition to commercial production. FEI (FEIC) builds a Focused Ion Beam (FIB) machine that uses a beam of gallium ions to perform stunningly fine repairs, etching, and deposition on completed devices. (FEI is also the leading manufacturer of complementary, ultra-precise measurement tools, such as dual-electron-beam and ion-beam microscopes.) Pulsed Laser Deposition (PLD) is already used widely in hundreds of industrial and academic R&D facilities around the world—carefully shaped ultrashort (nanoseconds) laser pulses ablate difficult materials (such as titanates), depositing ultra-pure atomic layers on target substrates. Pulsed Electron Beams (PEB) perform the same task, but allow higher deposition rates. Private and innovative Neocera has 80 percent of the U.S. market in PLD. (We were impressed enough by the company to invest in it ourselves earlier this year.)

Each of these tools depends on the very precise control of power to move atoms. In MBE, for example, extremely precise control of the temperature of the source controls the flux of material from each effusion cell. Very accurate control of layer thickness is achieved using precise temperature sensors (Tungsten-Rhenium thermocouples) and control loops. Chilled panels separate the cells to thermally isolate them from each other. The computer-controlled shutters in front of each effusion cell can turn off each beam within a fraction of a second. The growth of layers is monitored using high-energy electron diffraction.

Each of these methods has different strengths and weaknesses. Some grow fatter layers, less precisely but more quickly, others grow thinner layers, more slowly and accurately. Some use higher temperatures, which generally means they run faster, but are worse at maintaining a very clean divide across a junction. MBE is the slowest, the most difficult, and the most precise option of all.

Pure logic chips are built, for the most part, using the simpler tools. On these devices, all of the truly quantum engineering occurs at differently doped silicon-silicon interfaces. Insulators are built by oxidizing the silicon to make silicon oxide. Conducting wires are added in a relatively straightforward "metallization" phase. Most of the effort is directed at making structures smaller.

With optical devices and high-speed RF amplifiers, by contrast, the main challenge is to build the heterojunction interface itself. The quality and abruptness of the boundary is key. The power output, threshold current, and lifetime of solid-state lasers, for example, all depend strongly on these factors. CVD using metallic organic precursor gases (MOCVD) has emerged as the main engineering tool for building such junctions. But it can only be used when suitable precursor gases can be found, and maintaining interface uniformity is very difficult. Source gases are also highly toxic. So, for the most demanding and cutting-edge construction jobs, the quantum engineer turns to MBE.

MBE

The importance of the MBE machine can only be gauged by reference to what it is used to build. Motorola is already racing to commercialize the first big opportunity, the GaAs-on-Si wafer for high-speed radio-frequency amplifiers. Before it stumbled upon the possibility of depositing GaAs on Silicon, Motorola was looking for a way to deposit a better (higher dielectric) resistor. All the major manufactures of logic chips recognize that they will have to move beyond silicon-oxides if they are to build a billion-transistor CPU. If they end up using titanates, there's an excellent chance that Veeco's machines will end up as key tools in the logic-on-silicon arena, too. High-speed microprocessors also require power caches very close to the final loads. Microscopic aluminum and silicon-oxide capacitors are already stretched to their functional limits. Exotic titanate and metal versions built directly on the silicon will have to come next. MBE and IBD tools will be used to build them.

MOCVD remains faster and more economical for the deposition of the relatively fat layers used in light-emitting diodes, but MBE is now competitive on speed and cost for more complex, thinner layers, and superior in terms of precision, uniformity, and purity. MOCVD is good enough for building today's GaAs RF chips, and tomorrow's commodity LEDs; MBE machines will build the next generation indium phosphide chips (*The Power of Millimeter Waves, November, 2001*). MBE can handle both carbon and nitrogen as dopants; MOCVD can't. MOCVD can build a single-color light-emitting diode, but it takes MBE to build a true white-light device, consisting of a zinc-selenide substrate that emits in both red and green, with a doped (magnesium) ZnSe epitaxy that emits in the blue; the (tuned) combination produces white output without a phosphor. The best that MOCVD can deliver is a blue LED covered by a white-light-emitting phosphor.

As to cost, MOCVD has familiarity and history behind it, and the RF and opto-electronic industries have been built up around devices that tolerate its lower precision. MOCVD machines remain cheaper in fabs that already have the infrastructure in place to handle the highly toxic arsenine and phosphine precursor gases. But for the complex and precision junction devices, MBE is already cheaper than MOCVD machines that are being built from scratch.

The commercial MOCVD business is relatively mature. The leader in pure MOCVD is Germany's Aixtron. EMCORE Corporation (EMKR) is another major builder of commercial MOCVD machines. Nippon-Sanso is another player; so too is Novellus Systems (NVLS), a company that builds both CVD and PVD (Physical Vapor Deposition) machines, along with other more conventional wafer fabrication systems. One finds as well innovative (private) niche players, such as Structured Materials Industries started by a former EMCORE executive. A number of other semiconductor equipment vendors, including Britain's DEVP, Mattson Technology (MTSN), and CVD Equipment, build various other CVD machines.

Commercial-scale MBE machines, by contrast, are still at a much earlier stage of development. Veeco's MBE machines have larger capacities than its competitors' and can hold larger wafers. It is the only company with a commercial MBE machine that holds seven 6-inch-diameter wafers, with 4-by-8 and 2-by-12 machines coming soon. MBE rival Riber (part of Instruments S.A.) has announced a 7-by-6 machine. The serious competitor among the handful of other MBE tool-makers is Britain's VG Semicon, now owned by a U.S. parent, ThermoElectron (TMO). Other players, such as the U.K.'s Oxford Applied Research, provide only MBE sources, not the entire tool.

With Applied Epi in its fold, Veeco would have a prosperous business if it merely continued to build one or two of these enormously complex quantum tools. But the company is implementing a much more ambitious strategy—to acquire, develop, and integrate all of the critical tools required to build truly integrated circuits on the quantum assembly line of the future.

Veeco's Ultra-Clean Assembly Line

The inexorable tendency of the universe is to mix atoms up. Nothing short of utter emptiness can stop that process. Intel ranks its "clean rooms" in terms of tens or hundreds of dust particles per cubic foot of air. But in the seriously clean rooms, air is as unwelcome as dust. The assembly line of the future has to have one ten-trillionth the normal concentrations of any molecule, of any kind.

Veeco started out in 1945 as the "Vacuum Electronic Equipment Company." Today, Veeco is developing the fully integrated quantum machine shop—one that can perform a full range of interface-assembling functions, with assembly-line efficiency, in perfectly empty space. Veeco builds a manufacturing environment that maintains pressures a trillion times below atmospheric. And it integrates into that environment an array of tools that apply power to move material to build atomic-scale structures with assembly line efficiency. Veeco's systems are engineered, in other words, to reconcile a perfect vacuum with a high material throughput.

The six-foot-high, twelve-foot-wide collection of gleaming steel looks like it has been lifted straight out of a Star Wars set. The central vacuum chamber is three feet or more in diameter. Inside the chamber, wafers are mounted on a "planetary motion" platter that rotates on command to expose a particular wafer to a different machine tool, and to spin wafers as needed to promote more uniform deposition. The tools are mounted around the circumference of the chamber, directing material through discrete ports. A layer of material is deposited or removed; the platter is rotated; and another layer is deposited or removed.

Through acquisitions and internal development, Veeco has developed the capacity to manufacture almost all of the individual tools that constitute a complete cluster machine. Working with a private laser manufacturer, the innovative Nova Crystals, Veeco recently developed a unique reactive ion-beam etching tool, which provides the extremely precise deep etching capability needed to build edge-emitting lasers on a large substrate. A 1997 acquisition gave Veeco its first physical-vapor-deposition tools, along with the planetary tool-set for holding, moving, and spinning wafers. A 1999 acquisition brought in ion-beam etching (IBE). A year 2000 acquisition brought in additional PVD capabilities, along with atomic layer deposition (ALD) and MOCVD tools.

In various other transactions completed in the last

Quantum Machine Shop

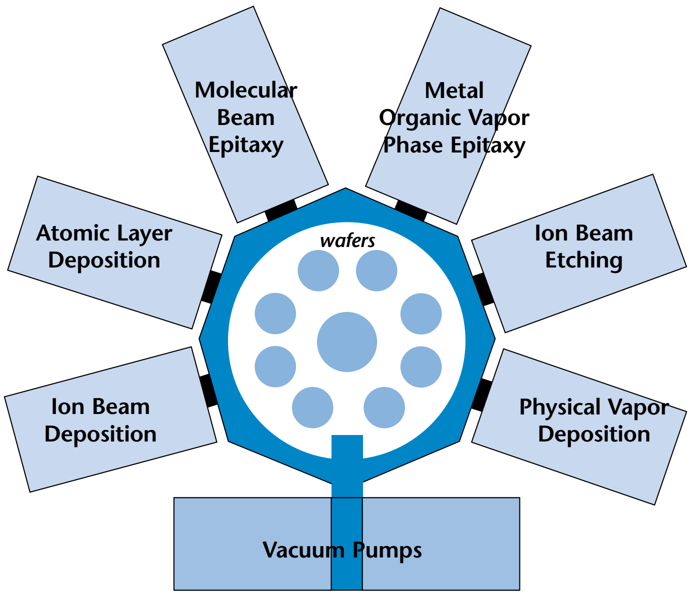

At present, each Veeco cluster machine is configured for a specific device—a laser, RF amplifier, or magnetic memory—with the combination of tools dictated by the kinds of quantum layers needed. The emerging ultimate cluster-tool will deposit a wide array of compound atomic-scale layers with assembly-line efficiency. Nobody yet builds a single chip that integrates logic, memory, lasers, and sensors. But such chips will get built. Integrate all these functions on the same silicon substrate as the digital logic, and you have a truly integrated circuit. Call it the integrated-integrated circuit, or I^{2}C.

five years, Veeco has also picked up manufacturers of high-current-density ion sources, atomic-force and scanning-probe microscopes, and laser interferometers, along with capabilities in quasi-static testing, and automated optical-defect inspection. The company is currently forging a partnership with Epion, a tool-maker currently owned by JDS Uniphase, to adapt their Gas Cluster Ion Beam process for use in the manufacture of next-generation magnetic materials. (This process accelerates a cluster of ionized argon atoms to produce ultra-flat surfaces by literally hammering atomic-level surface crystal defects back into place.)

Last year's purchase of Applied Epi added MBE to Veeco's tool cabinet. Veeco was already a leading supplier of ion and gas sources for the other suite of quantum tools. Now, with Applied Epi in hand, it has become the leading provider of production-scale MBE machines; it is also now a major supplier of effusion sources for MBE machines built by competitors. Applied Epi itself was founded to lead development of MBE sources, but the company moved aggressively into full-scope tool production after it acquired R&D MBE systems from PerkinElmer and Varian.

Veeco builds bigger vacuum chambers and larger effusion sources. By adding effusion sources and increasing the distance from the source to the target, Applied EPI has developed MBE machines that are capable of building layers over much larger areas than used to be possible. Veeco has also developed ways to manufacture very long-lived primary sources. (Applied Epi excels in this area, too—it is the largest supplier of effusion sources used on competitors' MBE machines.) These strengths make Veeco's MBE machines adaptable to higher volume production. And by clustering tools around a single evacuated core, Veeco machines allow multiple stages of manufacturing to occur in what amounts to an assembly line process within the chamber itself. Once wafer crystals and the source materials for epitaxial layers are loaded, and the machine is locked and evacuated, a continuous manufacturing process begins, and it can continue until construction is completed.

At present, each Veeco cluster machine is configured for the specific device being built—a laser, RF amplifier, or magnetic memory, for example. The choice of tools and sources is dictated by the kinds of quantum layers that the machine will build. For example, one market that accounts for 25 percent of sales of Veeco equipment is the manufacture of magnetic and dielectric surfaces for disk drive systems. These layers involve a "stack" of up to 9 layers, each one from about 10 to 150 Angstroms (A) thick. (By comparison, Intel's smallest gates are 1300 Angstroms.) The Veeco cluster that builds these layers includes two primary drivers (physical vapor deposition and ion-beam deposition) along with aluminum, cobalt, iron, and tantalum sources. PVD deposits the comparatively thick (100 to 200 A) cobalt, iron, and nickel layers. PVD is also used to deposit elemental aluminum in much thinner (10 to 15A) layers, but a separate ion beam tool is then used to fire oxygen atoms at the aluminum layer to oxidize it. Atomic Layer Deposition is then used to lay down sequential "monolayers" (one atom thick) of aluminum oxide for the ultra-high-precision electrical contacts. And Veeco is adapting its Gas Cluster Ion Beam tool to achieve the perfect (atomic-level) surface flatness required for perfect magnetic field uniformity.

Magnetoresistive Random Access Memory (MRAM) can now be built directly on silicon, too, using the same magnetic metals, insulators, and connectors. This delivers the high speeds of chip-based dynamic memory, at densities and costs more typical of spinning hard drives—and the memory isn't volatile, it doesn't forget everything when the power is turned off. Motorola announced in February a commercially viable architecture for MRAM on silicon; IBM, Infineon, and others, are also moving rapidly toward commercialization.

Veeco's ambition, however, extends far beyond any single application. The company is aiming to build the ultimate cluster-tool platform, that can build compound atomic-scale layers with assembly-line efficiency. The manager of the new evacuated factory will be able to choreograph the sequential engagement of every major quantum tool in the cabinet. By its own account, Veeco aspires to be the "first equipment company to integrate silicon manufacturing concepts and models into the compound semiconductor industry."

Feynman's Microcosm

"What would happen if we could arrange the atoms one by one the way we want them?"

"What could we do with layered structures with just the right layers?"

Richard Feynman asked these questions in 1959, and sketched out the answers to them, too. He described why we needed extremely tiny computing elements, and foresaw that such elements would be made at scales of "10 to 100 atoms in diameter." He foresaw solid-state light: "It is possible, for example, to emit light from a whole set of antennas, like we emit radio waves." Feynman, the physicist, even took a stab at sketching out how the engineers might proceed to build things in the "staggeringly small world that is below." He suggested that they might use a "source of ions, sent through the [electron] microscope lenses in reverse, [that] could be focused on a very small spot." Another possibility for "putting atoms down in a certain arrangement, would be to evaporate the material."

The only thing that Feynman didn't fully grasp is how difficult this kind of engineering really is. "In the year 2000, when they look back at this age, they will wonder why it was not until the year 1960 that anybody began serious work to move in this direction," Feynman declared. He was wrong about that. No one has to wonder much why Veeco didn't happen forty years earlier. Physicists—even great ones, like Feynman—always underestimate the practical engineering challenges, the nitty gritty obstacles that can be brushed aside in theory, but that make all the difference in practice.

And even engineers often underestimate how much easier it is to build a single device than to get it into mass production. The very idea of mass-producing quantum structures of atomic-scale precision almost defies belief. Mass production is the realm of high-speed assembly lines moving tons of material. Quantum engineering is the realm of painstakingly slow deposition of nanograms of material. Factories count their output in millions of life-sized units, devices that can be packed into boxes or mounted on circuit boards. Quantum engineers count their output in hundred-atom layers assembled in near-perfect vacuums.

Yet this gulf can be bridged. It has been in the manufacture of integrated logic circuits. Applied Materials (AMAT) was founded in 1967 with just five employees. At that time, most of the infant semiconductor companies built their own fabrication equipment. Applied Materials started supplying them with materials and components. Before long, however, Applied Materials was offering turnkey chip-making systems. In 1972, with sales of $6.3 million and 155 employees, Applied Materials went public. The company went on to develop precision etching machines, followed by the "Precision 5000," the first single-wafer, multi-chamber processing systems. The innovative architecture and automated handling capabilities of this machine revolutionized semiconductor manufacturing. The first unit built is now part of the permanent collection of the Smithsonian Institution. Applied Materials had $500 million in sales in 1989; ten years later it hit $5 billion.

The gulf between quantum engineering and mass production has been bridged a second time, in the manufacture of solid-state lights, lasers, RF amplifiers, magnetic media, and a wide range of other sensors. Effective processes, MOCVD prominent among them, have been developed to manufacture large numbers of these heterojunction devices at low cost.

Now the challenge is to get from IC to I^2C. As Motorola has recognized, there is an enormous payoff in just finding a way to build gallium arsenide amplifiers on silicon wafers, rather than on a GaAS substrate. The next step will be to put logic, and memory, on the same substrate, to combine logic and photons, or thought and light—hence Motorola's "Thoughtbeam" venture. Scientist Jamal Ramdani used research-level MBE equipment to build the first GaAs-on-Si device at Motorola's Tempe Research Facility. Now, the first order of business for Thoughtbeam is to get the process set up to put GaAs devices on 12" silicon wafers, at production-line volumes. Veeco's technology VP, Hari Hegde is Ramdani's engineering counterpart, the man who converts the scientist's instrument into mainstream production tools. When we talked to Hegde he conveyed the measured enthusiasm of a man who knows he's riding the next great technology wave.

RF, optical, memory, and sensing capabilities will inevitably end up integrated with silicon logic because

The Power Panel

Ascendant Technology	Company (Symbol)	Reference Date	Reference Price	6/28/02 Price	52wk Range	Market Cap
System Integrators	Veeco Instruments (VECO)	6/28/02	23.11	23.11	19.90 - 41.85	671.2m
	Oceaneering Intl (OII)	5/31/02	31.01	27.00	13.96 - 32.17	664.0m
	Amkor Technology (AMKR)	4/2/02	21.85	6.22	3.62 - 24.79	1.0b
	Emerson (EMR)	5/31/00	59.00	53.51	44.04 - 66.09	22.5b
	Power-One (PWER)	4/28/00	22.75	6.22	5.26 - 17.40	492.0m
Electron Storage & Ride-Through	Kemet Corp. (KEM)	5/1/02	19.63	17.86	13.85 - 22.40	1.6b
	Wilson Greatbatch Technologies (GB)	3/4/02	25.36	25.48	21.20 - 39.00	533.4m
	C&D Technologies (CHP)	6/29/01	31.00	18.02	16.35 - 33.90	468.2m
	Maxwell Technologies (MXWL)	2/23/01	16.72	8.72	5.81 - 22.50	93.8m
	American Superconductor (AMSC)	9/30/99	15.38	5.46	3.85 - 27.24	112.1m
Project, Sense, and Control	Danaher Corp. (DHR)	1/29/02	61.56	66.35	43.90 - 75.46	10.0b
	FLIR Systems (FLIR)	1/9/02	41.64	41.97	17.20 - 59.50	701.7m
	Analogic (ALOG)	11/30/01	36.88	49.17	33.40 - 56.50	651.9m
	TRW Inc. (TRW)	10/24/01	33.21	56.98	27.43 - 57.05	7.3b
	Raytheon Co. (RTN)	9/16/01*	24.85	40.75	23.95 - 45.70	16.2b
	Rockwell Automation (ROK)	8/29/01	16.22	19.98	11.78 - 38.60	3.7b
	Analog Devices (ADI)	7/27/01	47.00	29.70	26.60 - 52.74	10.9b
	Coherent (COHR)	5/31/01	35.50	29.99	25.05 - 39.50	866.0m
Powerchips	Cree Inc. (CREE)	4/30/01	21.53	13.23	10.35 - 33.32	963.1m
	Microsemi (MSCC)	3/30/01	14.00	6.60	6.95 - 40.10	190.6m
	Fairchild Semiconductor (FCS)	1/22/01	17.69	24.30	13.76 - 32.03	2.8b
	Infineon (IFX)	11/27/00	43.75	15.49	10.71 - 27.55	10.7b
	Advanced Power (APTI)	8/7/00	15.00	14.49	6.50 - 15.25	150.1m
	IXYS (SYXI)	3/31/00	6.78	5.38	4.27 - 16.25	170.2m
	International Rectifier (IRF)	3/31/00	38.13	29.15	24.05 - 50.50	1.9b

Note: This table lists technologies in the Digital Power Paradigm, and representative companies that possess the ascendant technologies. But by no means are the technologies exclusive to these companies. In keeping with our objective of providing a technology strategy report, companies appear on this list only for the core competencies, without any judgment of market price or timing. Reference Price is a company's closing stock price on the Reference Date, the date on which the Power Panel was generated for the Digital Power Report in which the company was added to the Table. All "current" stock prices and new Reference Prices/Dates are based on the closing price for the last trading day prior to publication. IPO reference dates, however, are the day of the IPO. Though the Reference Price/Date is of necessity prior to final editorial, printing and distribution of the Digital Power Report, no notice of company changes is given prior to publication. Huber and Mills may hold positions in companies discussed in this newsletter or listed on the panel, and may provide technology assessment services for firms that have interests in the companies.

* *The October 2001 issue closed on September 16, 2001 and was posted at 8 a.m. on September 17, 2001. Due to the markets' close in the week after September 11, our reference price reflects Raytheon's closing price on September 10, 2001.*

integrating more components on a single chip shrinks things down drastically. That is why capacitors and resistors were first integrated on semiconductors four decades ago, and why "integrated circuits" are now so ubiquitous. Integration makes the entire circuit faster, more robust, and ultimately, much cheaper. Chip-scale engineering changes everything, at least once you learn how to do it on a mass-production assembly line.

To get there, you need the machine tools and integrated production facilities to do the integrating at mass production speeds. So long as heterojunction devices are manufactured separately, they can be manufactured at their own pace, then packaged separately, and wired up to silicon logic ICs somewhere far downstream of the fab. The LED, laser, and RF amplifier manufacturers can produce tens of thousands of units per wafer, but by comparison with the logic-chip (silicon CMOS) assembly lines, the small, static, batch-mode processes used to manufacture lasers are leisurely and cumbersome. The integrated assembly will have to get up to logic-chip production speeds for the I²C to become a practical reality. And that is where Veeco is taking things.

Peter Huber and Mark Mills
June 28, 2002

More information about digital power technologies is available on www.digitalpowerreport.com