

The Law of the Powercosm: Burn Silicon

As powerchips ascend to a \$100 billion business, the stage is set for Intels and AMDs of the Powercosm to emerge.

If you're looking for 20 percent margins—on mileage, pollution, ROI, whatever—buy a Honda. Most power analysts are looking in that general direction, focused on yesterday's problems and opportunities. Theirs is the eighteenth-century paradigm of Sadi

Carnot, the man who first worked out the theoretical limits of thermal engines and the second law of thermodynamics. They measure "progress" by tiny upticks in thermodynamic efficiency. Or else they look to the paradigm of Thomas Malthus, Carnot's contemporary. The paradigm is centered in the dismal world of the green regulatory state, where "progress" is measured by even tinier downticks in emissions of trace pollutants.

There is nothing much wrong with efficiency or clean air, except that all the opportunities for delivering more of the same are plodding and small, and getting smaller. Carnot had reason to worry about efficiency: His steam engines were 5 five percent efficient. But we're at 70 percent today. Engines have likewise grown a hundred times less polluting in the last two decades alone. The old paradigms are just that: old, weary, exhausted. They address yesterday's receding scarcities. The new power paradigm addresses the overarching new scarcity: a three-orders-of-magnitude shortfall in the quality of electric power, and the rich new opportunities to meet that scarcity with rapidly evolving new technology.

Clean electricity is the fuel of the digital infrastructure. The market is already paying hundreds to thousands of times more for ultra-clean "six 9s" power than for old-fashioned two or three 9s. Those are the kinds of mark-ups we already pay for the electrons delivered by an "uninterruptible power supply" (UPS), whether at the desktop level from American Power Conversion, or at the telecom-hub level from Liebert (an Emerson Electric division). The value of electrons is now tied to the opportunity costs of doing without them. The overarching new scarcity: high-9s electric power.

The emerging new abundance—the technological key to overcoming that scarcity—is the power semiconductor. Solid-state power switches have been around for decades—they antedate even the integrated circuit. But until recently, they just weren't very important. Lower power powerchips were sold mainly into mature or declining markets, switching a few watts to a few hundred watts in TVs, stereos, PCs, as well as electric motor controls in industrial and commercial equipment. Expensive powerchips at the ultra-high-power level were sold, in very limited numbers to the heavy-iron power market, electric utilities. It is in the middle of these two poles that the rich new territory of powerchips is to be found, tightly tied to the ascendant technologies of the telecom.

The \$100 Billion Powerchip

Recent improvements in materials, device architecture, and manufacturing infrastructure have brought the powerchip to the golden knee of the technology curve, the sharp bend in the hockey stick of growth—about where microprocessors stood around 1980. A \$12 billion global market today, power semiconductors will be at least \$100 billion by the end of the decade. There is burgeoning new demand for devices to switch kilowatts and megawatts for servers, routers, wireless base stations, digital factories, dot.coms, micro-power networks, and distributed power storage and generation. A huge,

rapidly evolving, and still largely untapped market for silicon power switches is now opening up in the vast, untapped reaches higher up the power curve.

The Silicon Supply Side

Serendipitously, a simultaneous revolution is now under way on the supply side. The infrastructure to mass produce new, higher-power powerchips is now being propelled by technologies developed to manufacture memory chips, microprocessors, and integrated circuits. The key raw material of the powerchip? Silicon. The basic architecture: The silicon wafer. The powerchip industry is appropriating the materials and manufacturing technologies that the smartchip industry developed and continues to improve: the aligners, photolithography machines, and other equipment used to cut, handle, and etch the billions of wafers that pave the microcosm. Powerchip manufacturers that efficiently adapt the technologies of the smartchip fab line to powerchip manufacturing will gain an overwhelming competitive advantage.

The winning powerchip companies will push not very exceptional powerchips down the precipitous slope of automated mass production. They won't use the most exotic materials, or build the largest wafers, or the highest-power assemblies, compared with those to be found at the outer edges of their industry. They will handle too little power to be of much interest to the established "mainframes" of the power industry, large utilities. Little noticed by those who own and control the top of the power grid, they will redefine the infrastructure of electric power from the bottom. From the bottom up, they will drag the largest and most important engineering structure of the twentieth century into the twenty-first. They will enable and impel the massive transition from analog to digital power.

The Intels and AMDs of the powerchip industry are ready to emerge. But who are they?

Digital Power

The wonder of the logic gate was that it transformed the world of Edison into the world of Grove and Gates—raw power into machine intelligence. Now, the logic gate moves back into the world of power.

Bits are electrons; petabits of traffic now move through the telecosm; and gigawatts of electric power are now required to create, move, and store them. But not ordinary electrons. The telecosm demands unusually clean, stable, reliable electrons—"high-9s" power, available

clean and steady at least 99.9999% of the time. From wireless base stations, to routers, to factory-floor thin-servers, to local caching systems and data centers, the building blocks of the information economy demand at least six 9s, which is to say, at least three orders of magnitude better than the old grid currently supplies. The newest generation of nodes of the telecosm are the terabyte-level data local storage centers, now multiplying at an 80 percent annual rate. When it comes to these so-called storage area networks, "people think the game is in the memory," says Nathan Zommer, CEO of IXYS, one of the powerchip companies at the epicenter of the silicon powercosm. "But the limiting factor for reliability is actually in the power management" for the kilowatts and megawatts that a terabyte center requires.

The new metrics of quality extend from the level of the motherboard to the building and beyond, up into the networks, from the nanowatt levels of power amplifiers in cell phones to the gigawatt power levels in the backbone of the grid, an eighteen-order-of-magnitude span in power range. While the digital economy remained physically confined to stand-alone desktops and a limited number of mainframes, it consumed only a small fraction of our electrons, and the business of cleaning up power and achieving always-on reliability remained correspondingly narrow. But computers and the hardware of the telecosm now consume what we estimate to be 13 percent of our electrons. And their share will grow fast, as networked microprocessors come to permeate everything. The demand is no longer for isolated islands of smart—i.e., clean, uninterruptible—power, but is for smart power everywhere, all the time. Businesses already need power of that quality across the entire enterprise, not just at a scattered selection of individual desktops. A steadily rising fraction of well-wired residential consumers will demand equally good power household-wide. So will every node in the metastasizing wireless infrastructure.

As things now stand, no more than 15 percent of the electric power moving through the lower-power tiers of the grid, and virtually none in the high-power tiers, is conditioned or controlled by silicon. But powerchips are now poised to take over control of the entire electric power network, from the home appliance up to the central power station, the full length of electricity's own World Wide Web. The key to manipulating, shaping, storing or moving electrons at high-9s levels of reliability is the switch—by and large, the silicon switch—installed at all power levels of the electric network and controlled by software.

Editors	Peter W. Huber Mark P. Mills
Publisher	Richard Vigilante
Managing Editor	David S. Dortman
Design	Julie Ward
President	Mark T. Ziebarth
Chairman	George Gilder

The Digital Power Report is published monthly by Gilder Publishing, LLC. Editorial and Business address: P.O. Box 660, Monument Mills, Housatonic, MA 01236. Copyright 2000, Gilder Publishing, LLC. Editorial inquiries can be sent to: peterhuber@powercosm.com or markmills@powercosm.com. Single-issue price: \$50. For subscription information, call 800.261.5309, e-mail us at help@powercosm.com, or visit our website at www.powercosm.com

Power switches are already ubiquitous, of course. Deployed throughout the trillion-dollar electric power grid, switches determine where current flows, and how reliably. There are hundreds of millions of them, from the little wall toggles in your house that switch 100W light bulbs to 20 foot circuit breakers that switch megawatts of power at the doorstep of central coal and nuclear power plants. But they're the wrong kinds of switches. Almost all these switches are electromechanical, spring-loaded levers. The technology has advanced little beyond that which Thomas Edison used in his first Pearl Street Station power plant in New York City in 1882. Like the telecom grid of 30 years ago, the power grid is almost entirely analog. Its switches depend on Newtonian technologies that operate at the inherently slow speeds of the macrocosm.

That makes it far too dirty and unreliable for computers, the digital home and factory, the dot.com companies, or wired and wireless telecom providers. Electromechanical switches are slow, and slow switches are inherently stupid—when in action, they spend much of the time that counts in the ambiguous state halfway between “on” and “off.” As their metal contacts close, the current jumps out ahead of the metal. The resulting arc—a small bolt of lightning—creates electrical noise. For a brief period, and over an (often) considerable distance upstream and down on the grid, a cacophony of electrical spikes and pulses dirties the theoretically smooth 60 Hz sine wave of power, threatening equipment anywhere in its path. A system of messy analog switches creates a continuous, unpredictable array of states between on and off: dirty power that crashes digital devices. Clean, “digital” power requires switches that operate at 1 kHz and up, substantially faster than the 60 Hz clock of the electrical mains.

The technology exists: it is the solid-state powerchip. A powerchip's individual gates flip in microseconds or nanoseconds; complete modules can operate at 5 kHz to several MHz—some three to seven orders of magnitude faster than an electromechanical switch. At such speeds, a switch appears to the grid as though it has only two, crisp states: open or closed. The noise, the dirt, the system-wide disturbances vanish.

And this is what permits fast power switches to squeeze new 9s of power quality out of redundant generating systems. A large part of the pursuit of six-9s power centers on distributed generation. Some 50,000 to 70,000 MW of 0.2 to 10 MW micro-units (“micro” by big-utility standards) have already been deployed for stand-by power purposes in factories and buildings across the country. (This in itself is an astonishing number—standby capacity is approaching the entire nuclear generating capacity of the United States, and it is growing very fast—a bonanza for builders of micro-turbines, diesel gensets, and fuel cells.) But these back-

up systems are useless as providers of high-9s power unless they can be switched in cleanly at speeds unheard of for electromechanical switches. Fast switches can be synchronized “perfectly,” so far as the rest of the grid is concerned, and that makes possible clean, seamless load hand-offs between the grid and all lower tiers of backup generators, flywheels, and batteries. Fast switches also make possible clean load shedding—temporarily shutting off air conditioners, heaters, or refrigerators, for example—which, at the margin, is equivalent to adding backup capacity.

Recent improvements in materials, device architecture, and manufacturing infrastructure have brought the powerchip to the golden knee of the technology curve, the sharp bend in the hockey stick of growth.

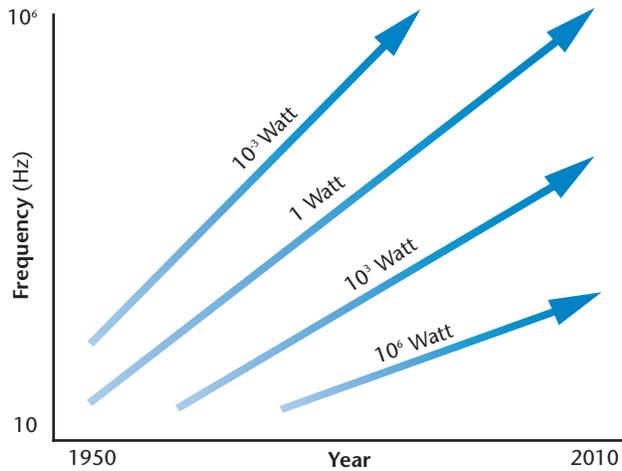
The powerchip, in short, defines a completely new paradigm of power switching. It doesn't improve on the electromechanical switch incrementally; it improves upon it by three to six orders of magnitude, along the single dimension that matters the most: raw speed. It makes possible, and propels, a fundamental, radical restructuring. The old model of the power grid was: Centralize and dispatch. The new will be: Disperse and switch. The old model: Chase thermodynamic efficiency by making fewer plants bigger. The new: Chase reliability by adding many more, distributed, power sources. The old model: Try to control everything, very intelligently, from a few critical control points. The new: Forget about perfecting the inherently dumb, inherently unreliable thermodynamics at the center—compensate with intelligent—i.e., fast—switches and distributed generation and storage around the edges.

Speed

Speed is the single, most critical dimension of powerchip performance. The faster the smartchip it stands behind, the more urgent the need for speed in the powerchip itself. For powerchips at the motherboard level, the rough rule of thumb is a powerchip speed about 1,000 times less than the clock speed of the microprocessor, according to Alex Lidow, CEO of International Rectifier. Each gigahertz processor in a state-of-the-art server requires its own power supply, with a megahertz powerchip inside.

Speed translates into power: The faster the chip, the more power it can handle. A chip's switching interval—the brief time it actually takes to open or close—is when it “stands alone,” exposed to the full current load that the upstream current pipe can punch through it. If the powerchip is going to melt, this is when it will happen. Other things equal, the faster the switching speed,

Powerchips Hit Telescopic Speeds



the tougher the chip. Powerchips can also be arrayed in parallel or stacked in series. The faster they can be turned on and off, the more precisely they can be synchronized. With good synchronization, even quite unexceptional powerchips can perform almost any function, including switching high levels of power.

Indeed, power technology that performs at digital speeds ultimately allows complete transparency, and full software control. Any power source (AC or DC) of any load can be controlled by simply assembling a suitable array of plug-and-play powerchip modules. If the components are fast enough, powerchip systems are fully scaleable. The rise of the powerchip will unleash a torrent of new development in complementary products, among them powerchip assemblies and peripherals, software to control powerchip arrays, software-controlled power filters, transformers, transfer switches, and other substitutes for the existing analog mainstays of old-world electrical engineering.

Fast Switch, Small Switch

And finally, higher speed translates into smaller size. Slow switches have to be surrounded with arrays of capacitors, inductors, and other analog components required to clean up the electrical noise the switches themselves create. Elaborate mechanical arrangements—mega-volt and kilo-amp inductors and dielectric liquids and solids, elaborate cooling systems, and the massive physical infrastructure to support them—are needed to damp out the disturbances. Fast switches don't make electrical messes. They can thus dispense with much conventional, power-electronic baggage. Improving a powerchip from 10 kHz to 50 kHz in a 250 kW device eliminates a half-ton of peripheral electronics and hardware, primarily the filters that are otherwise required to smooth the output of a system.

The smaller the powerswitch, the more widely it will be deployed. Base stations in the rapidly expanding

wireless market, for example, can require tens of kilowatts of power—which must often be squeezed into confined and expensive niches in church steeples, at the top of radio towers, or in high-rent urban buildings. The market wants a power-supply system with an energy density in the 100 Watt-per-cubic-inch range: That's like asking for a Corvette engine's power in a silicon breadbox. The size imperative is equally strong at the lower-power end, where today's second generation powerchip modules from firms like International Rectifier (IRF) sharply reduce component size and cost, and make possible redundant placement of power supplies in telecom racks, again raising the power 9s by replicating and dispersing the power assemblies.

Silicon IGBTs: In the "Low and Slow" Sweet Spot

The handling, routing, shaping, and conditioning of power electrons is accomplished across the entire power spectrum by a variety of device architectures, encompassing the usual acronymal alphabet soup: MOSFETs, FREDs, SCRs, GCTs, GTOs. But it is the insulated-gate bipolar transistor (IGBT), that is of greatest interest.

The IGBT typically consists of a single 6" wafer with several hundred individual dies; on each die are etched hundreds, often thousands of tiny MOSFETs driving tiny PNP bipolar transistors. Tiny by macroscale standards, but these are leviathan devices in the microcosm where tens of millions of transistors per die are the smartchip norm.

Like all solid-state power switches, IGBTs are much faster than electromechanical ones. But at speeds of up to 100 kHz to several MHz they are slow compared to the powerchip speeds theoretically obtainable on exotic materials like silicon carbide or diamond. And compared to the powerchips that the electric utility industry has been pursuing for years—for switching transmission lines at the 200 MW and up level—the IGBT architecture handles much less power. Dozens or even hundreds of IGBTs must be assembled in parallel to handle such ultra-high-power levels.

As silicon devices go, IGBTs are thus relatively "low (power) and slow (speed)." Massive arrays of such chips would be required to handle any serious long-distance backbone power levels. For that reason, IGBTs have been of little interest to traditional utilities, whose thinking invariably gravitates toward multi-megawatt power devices—like the cutting edge 10,000 kW powerchips in the power stratosphere. But IGBTs are still almost infinitely faster than electromechanical switches, and handle vastly more power than the powerchips already ubiquitous in such things as stereos and TVs.

IGBTs are thus positioned to bring digital power to the vast middle ground of the power market—entire homes, office buildings, dot.coms, wireless base stations, and storage area networks. The emerging wireless market

will require hundreds of thousands of transmitters on towers and telephone poles, each requiring high-9s power in the 10 to 50 kW power range. Millions of buildings and devices likewise operate at the 10 to 50 kW level; the lion's share of the entire commercial sector is in the 20 to 100 kW range. The silicon IGBT will be key to switching, controlling, and dispatching power from the vast array of generating systems now being built and deployed in the 20 to 100 kW size range: microturbines, diesel gensets, and fuel cells—all the various systems used to deliver short-wire, high-9s, backup power. Silicon IGBTs are also the powerchips that will find their way into large industrial motors—and thus into everything from industrial robots and lasers to electric welding systems and refinery pumps.

In sum, silicon IGBTs occupy the sweet spot in performance characteristics: about 20 kHz to 300 kHz and 10 to 100 kW. They are ideally suited for medium-power applications, and are poised to coast down the steep slope of silicon-centered mass production, and up the learning slope of faster and faster speed.

So who will make them?

In the Powerchip Sweet Spot

Dynex Semiconductor (www.dynexsemi.com) is a pure play in the powerchip market, a UK-based firm (trades on the CDNX as DNX) with extensive low-end powerchip capabilities. Intersil (www.intersil.com), formerly privately held by a venture capital arm of Citigroup with the technology assets purchased from Harris Semiconductor, went public on February 25, 2000. Intersil's powerchip technology capabilities cover a wide band of the critical sectors, from wireless through satellite and defense markets. Another substantial low-end powerchip player is Infineon (www.infineon.com), spun off from Siemens in March 1999. Infineon's recent IPO marked the largest in German history. All three are in the powerchip 'space,' but all have their dominant products at the low end of the sweet spot (at least for now).

Privately held (with an IPO in the future, we believe) Advanced Power Technology (www.advanced-power.com) is another pure play that competes in this realm and well up into the critical power level. And deep into the high-power end of the powercosm, serving automotive and industrial markets up into the ultra-high-power utility-network grade powerchips, are a wide range of powerchips from Powerex (a GE-Mitsubishi joint venture, www.pwr.com).

The pure plays in the public domain are International Rectifier (at www.irf.com) and IXYS (SYXI, at www.ixys.com). Both fit our powerchip paradigm very closely. Both have strengths in the key middle-power, high-speed market. Both manufacture their powerchips on state-of-the-art facilities, free-riding on the manufacturing technologies of the smartchip industries.

International Rectifier is the venerable pioneer in solid state switches (customers include Nokia, Sony, Lucent, Ericsson, APC, Emerson, Intel, AMD, Ford, and Siemens). It went public in 1958 and has the experience, technology, and market credibility to push silicon up the power curve. With nearly three-fourths of its business in advanced powerchips, the company's CEO has been practicing (and preaching) the gospel of the powercosm for years, albeit in slightly different language and (it would appear) to a largely uncomprehending audience. "We need to view power systems as integral parts of the entire electronic design," declared Eric Lidow the company's founder, chairman, and father of current CEO, Alex Lidow. He said that three years ago. The investment community hardly noticed. It will.

International Rectifier obtained broad patent protection for key IGBT architectures in 1983, at the dawn of IGBT technology. Its early IGBTs switched at about 2 to 3 kHz. By 1993 they were at 25 kHz; by 1995, 50 kHz. In 1997 the company introduced its WARP Speed™ switch, which hits 100 to 150 kHz, triple the industry norm. Architectural improvements followed in 1999, which boosted performance of its multi-kW IGBTs another 20 percent. Alex Lidow expects the IGBTs to approach MHz speeds within the decade. Increases in speed, together with complementary improvements in architectures and materials, are pushing up the company's powerchip module power densities by 30 percent every 2 years. Reduced component counts and lower operating temperatures are doubling reliability (mean time between failure) every 2 years. Switching losses (waste heat) have dropped six-fold since 1992.

The telecom demands unusually clean, stable, reliable electrons—"high-9s" power, available clean and steady at least 99.9999% of the time.

IXYS is on the same trajectory (customers include Rockwell, ABB, Emerson, Still GmbH, Eurotherm UK, and Alpha Technology). Founded in 1983, with products and primary revenues anchored in the powercosm, IXYS manufactures IGBTs that followed the same trajectory starting at only a few kHz, reaching now into the high-speed stratosphere going beyond 100 kHz, with power capabilities up to 100 kW. IXYS just announced a new IGBT architecture, a Reverse blocking IGBT (RIGBT) that has the potential to eliminate up to two-thirds of the physical components in front-end power management. A standard IGBT can block power flow in only one direction; the RIGBT blocks in both, which makes it ideal for operating in an alternating current environment without first rectifying the power. IXYS has the patents; the company's CEO says he'll manufacture all he can, and license others to

manufacture more. He has strong positive relationships with all the major players for such licenses, including the ABB—which holds 40 percent of IXYS stock. IXYS CEO Nathan Zommer told us that the power side of the equation is so critical to his customers in the global wireless market that the technology is almost entirely driven by “power density and power performance, not cost” in the search for the power. IXYS already sees over one-third of its powerchip business going directly into the telecom.

International Rectifier and IXYS also serve significant markets in lower-power (MOSFET) powerchips (important, for example, for battery charging controls in high-9 back-up systems). The MOSFETs, too, are pushing even higher and faster up the speed curve. But IGBTs represent their companies largest untapped markets. Both companies are close to pure plays in powerchips generally, and as pure as the plays come in the medium- and high-power-level IGBTs.

Both companies will also certainly face significant challenges from the likes of Toshiba, ABB, Mitsubishi, and Powerex Inc. (the closely-held venture for Mitsubishi’s products in a joint agreement with GE)—companies that sit at the high end, both in power terms and corporate girth. Toshiba and Powerex are heavily involved in solid-state switches at the ultra-high power levels, too. Toshiba is one of the few suppliers of the massive 6” optically triggered multi-megawatt thyristors, which can individually handle tens of megawatts. The optical triggers increase precision and greatly reduce the number of parts that surround the switching module that houses the core powerchip, but thyristors remain, by our standards, slow devices.

ABB, a company deep and wide into every aspect of the powercosm, is also engaged in the powerchip business and is consolidating its position at the high power end of that market. ABB has just announced a \$70 million investment in a new factory in Lenzburg, Switzerland, to manufacture high-end IGBT powerchips. ABB has pushed IGBTs up into the megawatt power range as well. At the same time, ABB has just introduced a variant on the IGBT and standard thyristor, called an insulated-gate commutated thyristor (IGCT). The IGCT targets the heavy-iron market of multi-megawatt powerchips, but operates at much higher speed than conventional high-end thyristors. ABB, not traded in the United States, has coyly announced revolutionary products in the power management and transmission areas and plans a U.S. IPO this year. With the combination of its powerchip position and its deep involvement across many other layers of the power industry, ABB is poised to emerge as the Lucent of the powercosm.

None of these larger contenders is a pure play in powerchips, however. Most resemble IBM or Motorola: Their

powerchip divisions are substantial, but operate inside huge, diversified corporations. Those divisions are likely to prove much more profitable over the next five years than is commonly recognized, with potentially significant, positive impacts on performance of the parent company. But because they operate in so many other markets as well, these companies are, at best, limited and cautious powerchip investments. And ironically, the powerchip may present as much threat as opportunity to these established incumbents of electrical engineering. All of them have huge legacy revenue streams anchored in analog switches, filters, and transformers—all of which can be completely eclipsed by digital power technology. The companies may manage their way out of this dilemma, but doing so will require some wrenching changes.

Riding On Silicon on the Engineering Coat-Tails of the Microcosm

High demand for high-9s power does not, of course, explain why a small group of companies is now so perfectly positioned to meet it. What explains that is a single word—“silicon.” Silicon wafers—of the right architecture—are why a small handful of companies are now certain to emerge as the Intels and AMDs of digital power.

As we have noted before, the powerchip came first—it antedates the smartchip. Yet powerchips have not been widely deployed, except at the lowest power levels. The companies responsible for the higher-power tiers of the grid—electric utilities, mainly—did not choose to deploy primitive, electrically noisy, analog switches. But not because the public was clamoring for noisy, unreliable power. They did so because analog switches were far cheaper than the digital alternatives.

But design and engineering factors have converged, quite recently, to make powerchips the cheap and reliable alternative, not only at the lowest power levels found in stereos and TVs, but also on up the power curve, at the power levels found in entire buildings and dot.coms.

Consider, first, how not to take over the switching of higher-load electrical power—not even with silicon powerchips.

Several companies already manufacture powerchips capable of handling 5,000+ V and 3,000+ amps (15 megawatts and up)—which is to say, suitable for switching power plant and grid-level loads. Such chips currently represent a \$1 billion market. There are two main architectures: The silicon controlled rectifier (SCR), and the gate turn-off thyristor (GTO). Ultra-high-power SCRs were first built over fifteen years ago, on thick, fat wafers. They are currently being used to isolate and interconnect high-power transmission lines in about 50 grid-level interconnect points around the world. The first GTO prototype was installed in 1996, to control a 100 MW transmission line operated by the

Tennessee Valley Authority.

These devices, however, are extremely expensive to build. All other things equal, higher voltages require thicker wafers, and higher currents require larger wafer areas—up to 6", currently, and 900 to 1,500 microns thick. Wafers that thick won't fit the aligners, photolithography machines, and other equipment used to cut, handle, and etch the millions of wafers that underlie today's microprocessors and smartchips—those wafers run 600 microns and under. So the powerchips at the ultra-high-power levels have to be made, in effect, by hand.

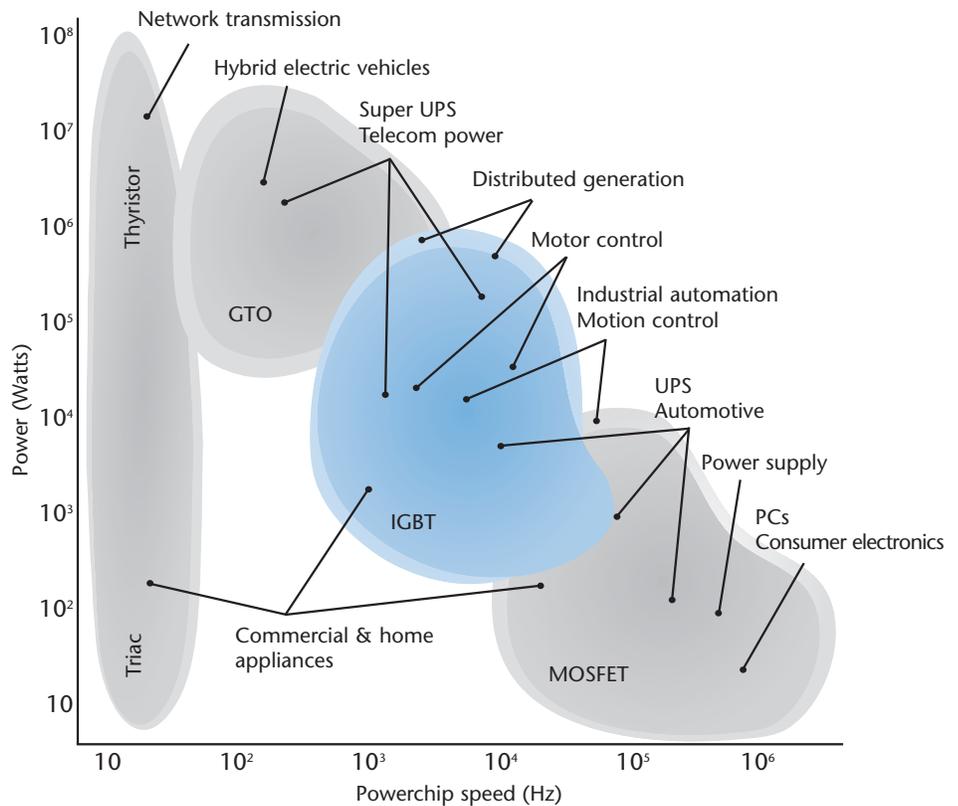
EPRI and DARPA are funding the pursuit of exotic materials and architectures that can push these high-power capabilities even higher. The future holds derivative powerchip materials such as silicon carbide (SiC), which are considerably faster than silicon, and run cooler. Cree (CREE) manufactures SiC for their famous blue LEDs, and has plans for the powerchip business; they could indeed emerge as a major force in the powercosm when SiC powerchips come of age. But that's the future, and there are enormous technical challenges still to be overcome before it arrives.

The digital world demands digital power today. The sweet spot for here and now is at lower power than the SCR and GTO monsters are painstakingly assembled to handle, and lower speed than the exotic next-generation materials may eventually deliver. And by happy circumstance, it centers on architectures, materials, and speeds that can be produced using the vast, versatile, constantly improving fabs already developed and widely deployed for manufacturing silicon smartchips. The powerchip companies that can ride down the Intel-AMD curve are those that are using the same old silicon—the most deeply studied material in history—and the same old manufacturing technology of smartchip silicon wafers—the manufacturing technology that has advanced the furthest and fastest of any manufacturing technology in history.

In due course, those same companies will bull their way up to the higher power levels too, but not with handmade devices. The right way to push powerchips up the power curve is the concept being pioneered by the Office of Naval Research (ONR). The ONR's Power Electronics Building Block (PEBB) program is focused on silicon and—above all—on developing a modular, plug-and-play

Powerchip Hierarchy

The mid-power market is where the explosive growth will happen in the Powercosm



design. The objective is to create standard power modules that can be stacked and assembled. Students of the history of the microcosm will recall that the integrated circuit itself originated in an analogous program, impelled by a vision of logic circuits that could be assembled out of plug-and-play electronic Lego blocks.

The Navy's interest in powerchips today stems from its interest in an all-electric ship that has superior reliability, dramatically reduced volume and weight, and a new generation of Buck Rogers-style electric-based weapons. The goal of the PEBB program: a 70 kHz, one megawatt (1,000 V, 1,000 A) plug-and-play device of mid-range performance and capability that can be stacked and assembled.

The powerchip architecture being pursued by ONR is the ultra-high-power MOS-controlled thyristor (MCT). It is fabricated on a single wafer with about 80 dies, each with 1 to 8 gates, etched on a single substrate. The MCT's simple gate structure can reach 25 kHz now, with 70 kHz in range. The MCT architecture's main advantage, however, isn't even its high speed, or the compact assemblies that speed makes possible. The main advantage is that, like its sister device the IGBT, it can be built with the same manufacturing technologies used to make smartchips. The "thin" MCTs now made

The Power Panel

Ascendant Technology	Company (Symbol)	Reference Date	Reference Price	3/31/00 Price	52wk Range	Market Cap	Customers
Powerchips: Insulated gate bipolar transistors (IGBTs)	IXYS (SYXI)	3/31/00	13 ⁹ / ₁₆	13 ⁹ / ₁₆	2 ¹ / ₂ - 21 ³ / ₁₆	\$166m	Rockwell, ABB, Emerson, Still GmbH Eurotherm Ltd. (UK), Alpha Technology
IGBTs	International Rectifier (IRF)	3/31/00	38 ¹ / ₈	38 ¹ / ₈	6 ¹ / ₂ - 48	\$2.33b	Nokia, Lucent, Ericsson, APC, Emerson, Intel, AMD, Ford, Siemens
Network Transmission and UPS: High-temperature superconductor	American Superconductor (AMSC)	9/30/99	15 ³ / ₈	44 ¹ / ₂	8 ³ / ₈ - 75 ¹ / ₈	\$851m	ABB, Edison (Italy), ST Microelectronics, Pirelli Cables, Detroit Edison, Electricite de France

Note: This table lists technologies in the Powercosm 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 judgement 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 of the month prior to Digital Power Report publication. 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.

on 600 micron silicon wafers (powerful enough to handle 6 kV and 2,000 amps at 25 kHz) can already be fabricated on existing smartchip fab lines.

The firm manufacturing powerchips for the Navy's PEBBs is Silicon Power (SPCO). It is privately held (with an IPO in the foreseeable future) and relatively small (\$15 million annual revenue), with a fab line in Pennsylvania and another in Latham, New York. Created by a former Electric Power Research Institute (EPRI) employee in 1994, SPCO acquired the high-end powerchip assets of Harris Corp. in 1998. SPCO also produces the completed PEBB modules and has logged over a dozen turnkey switch installations for high-9s power in the 5 to 38 kV range, for clients such as Ford, FMC, Texas Utilities and, TVA. The company has entered into an agreement with Rockwell to supply the U.S. DOE Partnership for a New Generation of Technologies automobile program with an MCT-based PEBB. The automobile angle is of limited direct commercial importance, but the three-year contract will likely lead to substantial advances in the 55 kW PEBB technology. And most important of all: SPCO not only covers the powercosm with their own line of IGBTs but, critically, they make their MCTs on a commercial IC fab line in San Jose that it rents from Micrel, and another line rented from a Sony fab plant in Texas.

As ONR and SPCO reach for the highest-power margins of the powerchip industry, however, IGBT companies like IRF and IXYS are already well established just one tier lower on the power curve. They have been in business for years, but now they are the direct, free-riding beneficiaries of extraordinary advances in silicon material processing and silicon wafer engineering. They did not have to develop the silicon foundries,

aligners, photolithography machines, and other equipment used to cut, handle, and etch millions of wafers at ever declining prices: the smartchip fabs did that for them. And they will continue to be the unintended but direct beneficiaries of every major advance in silicon science and silicon-wafer engineering.

They are seizing the opportunity. IXYS uses state-of-the-art hardware from SRAM smartchips, and, indeed, IXYS avoids the capital cost of a fab plant entirely—all its core products are produced on commercial SRAM and IC foundry lines in Korea and Japan. It gets better: SRAM lines are optimized to provide acceptable yields on fine sub-micron gates. But powerchips do not just tolerate but actually require larger gates (several microns to handle the power)—a breeze for SRAM production lines, and yields are very high. For its part, IRF uses a highly automated, state-of-the-art powerchip fab plant outside of San Diego. It runs a state-of-the-art fab line with full automation, but with equipment that has critical 'specs' a generation or so less demanding than those of the smartchip industry. Returns on capital are much higher (up to \$3 per \$1 invested) because of the massive free ride the company gets on wafer technology developed first for smartchip applications.

The law of the microcosm: Burn transistors. The new law of the powercosm: Burn silicon. The main challenge for powerchip companies is to take the triumphs already achieved in microcosm materials science and engineering and adapt them to their own rapidly emerging industry: the digital powercosm. That is exactly what they are doing.

Peter Huber and Mark Mills

April 14, 2000