THE HUBER MILLS digital power report
Powering the 21st Century
May 2002/Vol. 3 Issue 5
Published by Gilder Publishing, LLC

Tantalum, Titanates, and Silicon

Kemet's technology roadmap is focused on the leading edge of digital power and GHz logic

In January 2000, a kilogram of tantalum cost $65. By the end of the year it was at $550; triple the price of silver. (It has dropped back to the $60 range today.) What caused the spike? Surging sales of cell phones. And what does tantalum have to do with wireless telecom? Nothing, really. There's no tantalum in the DSPs, none in the ASICs, none in the memory chips, none in RF amplifiers, none in any of the I/O channels, or in the display. The tantalum is used to supply power. Tantalum capacitors store more power in less space than any other competitive alternative—which makes them essential to the delivery of highly ordered power to high-speed digital devices. The faster the digital logic, the more tantalum capacitors you need.

Moore's law observes that the number of gates that chip fabs could etch on to a given area of semiconductor surface doubles every 18 months. But nothing can double and redouble forever. Ever since Gordon Moore articulated it in 1965, pundits have wondered when his law would reach its inevitable limit. How finely can we project patterns on to silicon? How finely can we etch them into the semiconductor? How tiny can we make the ultra-thin, ultra-pure oxide layers that serve as insulators?

Despite chronic expressions of worry, however, chip engineers still have a pretty clear roadmap for how they'll keep things going forward from here—the 2 GHz 50-million-transistor Pentium—to the 100 GHz, and even the 1,000 GHz CPU and billion transistor chip. Photolithography has been pushed into the deep ultra-violet bands. Chip fabs have perfected precisely controlled "dry" plasmas and ions for etching. They've learned to deposit insulators in layers just dozens of atoms thick. They're at 130-nanometer lines in a gate today; they'll get to 20-nanometer lines and 100-GHz clock speeds in due course. 3-D architectures will push things further still, raising effective chip densities another order of magnitude, at least. A review of current literature from both Intel and the Semiconductor Industry Association (SIA) suggests that there are no barriers on this side of things that will prove insurmountable any time soon. The limit to Moore's law won't be reached in the logic gate at all. It will be reached in the flow of power into and out of the chip.

Most digital technology pundits see the power issues as peripheral and distracting, if they see them at all. Just assume suitable supplies of electrons, then get on with the serious business of building gates. But when you get right down to it, both the logic and the dynamic memory are nothing but "electrons," suitably channeled and confined. To say that digital logic "needs" power is to understate things badly. At the most fundamental level, the power and logic converge.

But not easily. This all-important interface depends on extremely challenging, and little understood technologies. As we discussed last month, this interface represents one of the most fecund opportunities in semiconductor-related technology (*Packing Power, April 2002*). It implicates packaging, voltage-regulating circuitry, electromagnetic radiation, and a host of thermal issues. In the middle of it all stands the venerable, much ignored, capacitor.

Headquartered in Greenville, South Carolina, Kemet Corporation (KEM) (http://www.kemet.com) is the world's largest manufacturer of solid tantalum capacitors, accounting for about one-quarter of that $2-billion market segment. It makes a wide range of other capacitors as well, including multilayer ceramic capacitors (made from titanates) that comprise an equally important $3-billion market segment, in which Kemet ranks as the second largest U.S. manufacturer and the fourth largest worldwide. Kemet is also a leading developer, one of only a few, in the emerging market for high-performance solid aluminum capacitors. Company wide, Kemet does capacitors, and only capacitors. Last year it shipped about 36 billion capacitors, from a portfolio of 35,000 different types. As of last month, it claimed to manufacture capacitors with the highest-power-density (per unit volume), together with the properties that make for the highest speed (i.e. the lowest "equivalent series resistance").

Kemet's sales are certainly down from the bubble days, when manufacturers of digital hardware were piling up huge inventories to build for growth that was never expected to end. Blank out the bubble, however, and you find that Kemet has followed a steady path of technology advancement and growth in the twelve years since the company's current management (with help from Citicorp Venture Capital) acquired Kemet's outstanding common stock from Union Carbide (1990), brought it public on the NASDAQ (1992) and then moved to the NYSE (1999). Semiconductors need capacitors, but the capacitors are in some respects a stronger and steadier market. Integrated circuit sales grew 11 percent annually during the past decade; capacitors have grown at a 20 percent annual rate during the same period. The worldwide market for all capacitors of all types and power levels now approaches $20 billion. There is also a marked shift in the capacitor market under way—toward the high-end, high-power-density, high-speed, surface-mount capacitors. The high-tech edge of the capacitor market is defined by a narrow class of materials, architectures, and applications that separate it from the commoditized rest.

Kemet's capacitors land in just about every form of electronic equipment, from commercial equipment and consumer appliances to military and aerospace systems. At Kemet's Web site you'll find no slick PR, no pretty pictures, no hyperbolic prose. The pages look like they come from the back pages of a company that supplies a thousand different nuts or bolts. The content is dense, cryptic—pure techno-geek impenetrable to much of the Wall Street and business press. The site is organized, instead, for the convenience of digital power's design engineers. Kemet's technology roadmap is focused single-mindedly on the leading edge of the demand: high-frequency, high-power-density capacitors.

Memory, Logic, Power

A capacitor is a bucket that holds electrons rather than water. Dynamic memory consists of banks of tiny capacitors—millions upon millions of them. Each capacitor, each individual bucket, can be full or empty, and a tiny sensor can tell which state it's in. You store a "1" by pouring charge into the bucket, and replace it with a "0" by draining it back out. Logic gates have an effective capacitance too. They're supposed to just open and close, but it takes some flow of current to make that happen, so you have to pump some charge in to flip the gate one way, and drain charge out to flip it the other.

Memory capacitors and gate capacitors are tiny, of course, and Moore's law says they'll keep getting smaller. Per bit, or per logic operation, the amount of charge required to store a bit or flip a gate keeps dropping and dropping. But the number of memory and logic capacitors keeps rising all the faster—now pushing into the *billions* on a single desktop computer. The smaller they get, the faster you can run them—that's at least half the reason for making them smaller in the first place. Gigahertz microprocessor speeds are now routine in desktop computers. The power amplifiers in wireless communications systems are now pushing up toward the tens of gigahertz. Millimeter wave radars are heading for the 100s (*The Power of Millimeter Waves, November 2001*). Power electronics in optical telecom systems now approach terahertz speeds. Comparable speeds (tens of MHz and beyond) are encountered in the electronics that drive optical displays, radar systems, and the best audio amplifiers.

Now consider what these speeds imply for power. The memory and logic capacitors are all getting smaller, but they're multiplying even faster—and how much power you need depends not on the individual memory bucket or gate, but on what they all add up to together—all 50 million transistors on a Pentium IV, for example. That typically means 10s to 100s of Watts of peak power, when the whole device is working flat out on the toughest challenge it's designed to handle. And the power must be conditioned and delivered as fast as the load itself can change—which means, roughly, as fast as the clock speed of the digital device itself.

Finally, these 10s to 100s of Watts, fluctuating at gigahertz speeds, have to be pumped into the device at close to 1 volt, and soon much lower. Logic devices *must* run at very low voltages, because the gates have been made so small, to make them so fast—any higher voltage will punch through the ultra-thin walls of the gates themselves. But to pump 100 Watts into a chip at 1 volt, you have to deliver 100 amps—an astounding amount of current for an industry that started out working in milliamps.

Authors	Peter W. Huber Mark P. Mills
Publisher	Lauryn Franzoni
Designer	Julie Ward
Subscriber Services	Melissa McNally
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.

Single-issue price: $50. **To subscribe call 800-223-0491**, e-mail us at hmpowerreport@gilder.com, or visit our website at www.digitalpowerreport.com

In the jargon of chip engineers, high current plus high speed translate into a truly gargantuan "di/dt," the parameter that measures how fast the current can change with time. Divide 100s of amps by billionths of a second and you have a challenge far bigger than any encountered elsewhere in the world of practical power engineering. It is the electrical equivalent of taking something that may not seem very big by interstate highway standards, but then requiring it to accelerate a billion times faster than a Porsche.

So how does one deliver power of this character? It's not at all easy. We briefly surveyed some of the reasons in last month's issue on packaging (*Packing Power, April 2002*). You can't begin to run serious current through ordinary circuits at gigahertz frequencies. At low speeds, a copper wire is (for the most part) merely a "conductor"—what goes in one end comes out the other more or less immediately, and that's that. But boost frequencies enough, and the wire becomes two other things—an inductor, and an antenna. Inductance is a measure of electrical inertia—push current in at one end of the wire and it doesn't just spill out the other, it piles up, and then begins sloshing back and forth in the wire itself. And when that happens, the wire also dumps power out in the form of radio waves.

There's no tricky way to dodge the problem. Use a fatter wire, and you get more inductance, not less. Use a thinner one, and you get less inductance but more resistance—so you dissipate less power as sloshing-current radio waves, but more power as heat. And the higher the rate of change—the higher the di/dt—the worse the problems get. The wire itself becomes the enemy. The only solution is shorter wires.

But you can't install a power supply in the Pentium itself. You can't even install one right beside the Pentium. The closest anything in the nature of a complete "power supply" can come is in the brick that's mounted on the motherboard of every server (*Cisco of the Powercosm, May 2000*). But the fastest bricks can't follow loads that vary much above 100 MHz. A very fast voltage converter has a reaction time of perhaps 100 nanoseconds. A microprocessor's current transients can change two to three orders of magnitude faster—over intervals of 1-20 nanoseconds.

In the end, ironically, the only way to power the most modern of electrical devices—the semiconductor logic and memory chip—is with the aggressive use of the very oldest of "electric power" technologies—the humble capacitor.

Power Density and Speed

Capacitors are built by the hundreds of billions and installed everywhere. No power supply, no electric system, can function without them. They range from the Mack-truck-sized oil-filled capacitors installed in grid-level distribution substations, to the dust-mite-sized low-power capacitors in every cell phone. They smooth the flow of power from megawatt stages of the grid down to the microwatt stages of digital power consumption, from the 700-kV aluminum wires that stretch across the countryside to the 1.5-V nanogrid that powers a Pentium.

No electric system can function without capacitors

A capacitor is an arbitrageur of electrons. It's the technology of choice for resolving short-term mismatches between power supply upstream, and load downstream. When supply suddenly dips, a capacitor can level it up, and when supply spikes, it can level it down. The capacitor is an old technology—about as old as any technology gets in the realm of electrical power—as old as Ben Franklin's kite, and older. Capacitor technology has already had a quarter of a millennium to mature. Yet today, it is developing faster than at any time before in its history.

We devoted an entire issue to capacitors a year ago, though in a very different context, the context of the hybrid-electric car. In the *Electron Cache* issue (*March 2001*) we discussed the kilofarad ultracapacitors manufactured by Maxwell Technologies (MXWL)—capacitors with a billion times more electron-storing capacity than those found adjacent to a CPU, capable of standing between slow-and-steady charging systems and up-and-down power requirements in diesel-electric buses and (eventually) hybrid cars. We returned to capacitors briefly in last month's issue, in our broader discussion of the power-centered technologies of chip packaging. A microprocessor (or an RF amplifier) on a circuit board comes surrounded by "decoupling" capacitors, mounted all around it like a small school of pilot fish. In some instances—as on the 72-Watt 600-MHz Alpha chip—some of the decoupling capacitors get built right on to the chip itself.

Decoupling capacitors store only a few cycles worth of charge at most; they must be constantly replenished from further up the line. But they provide just enough intermediate power caching to keep things hot even through the most demanding transients. The microfarad (μF) capacitors circle the Pentium to provide microseconds of load-leveling storage for the picofarad (millionth of a microfarad) capacitors that lie hidden inside all digital logic gates, and that define dynamic digital memory.

It takes quite an array of capacitors to do that. Flip a Pentium on to its back, like a beetle, and you expose not just the underlying ball grid array (the I/O) but also

Capacitors: Digital Power at Digital Speeds

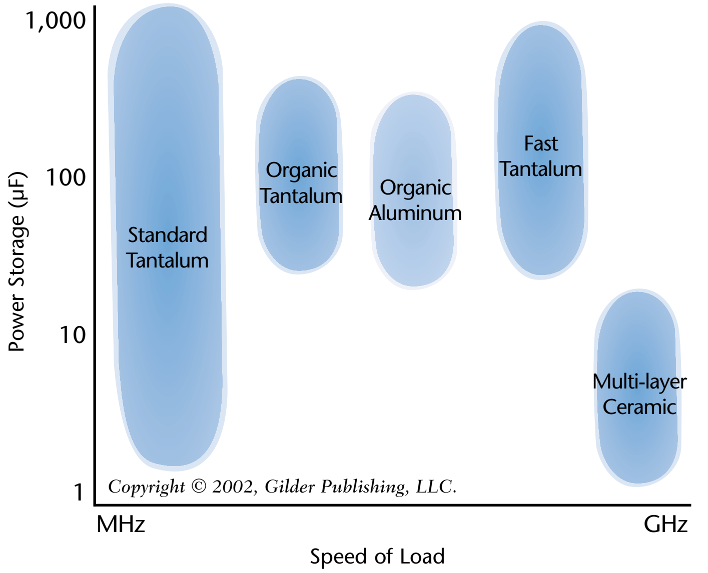

Digital loads require increasingly greater amounts of power at ever escalating speeds. No power supply can match the challenge without the intermediating use of a rising new class of very fast, very powerful capacitors. A capacitor is a bucket that holds electrons rather than water. Flip a Pentium on to its back, like a beetle, and you expose not just the underlying ball grid array (the I/O) but also attached leach-like is an array of both tiny ceramic and more powerful, but slightly slower, tantalum capacitors.

attached leach-like are a half-dozen or more tiny ceramic capacitors (from a fraction to a few µF each). A fraction of a centimeter away lies an array of 100-µF tantalum capacitors. Why so many separate capacitors? Why ceramics first? Why tantalum next?

Bigger capacitors store more electrons, but can't be mounted as close to the final load. So instead of raw size, the capacitor manufacturer pursues materials and architectures that store more charge in less space. *Power density*, in other words, is the first design objective. *Speed* is the second. The charge has to move in and out of the capacitor with little loss, at the speeds dictated by the ultimate payload. Higher speed requires less resistance to the passage of charge in and out of the capacitor. That quality is measured as "Equivalent Series Resistance"—the lower the ESR, the better—and "Equivalent Series Inductance" (ESL)—again, the lower the better. Both ESR and ESL change with the frequency at which you're trying to drive the capacitor—ESR, for example, drops as frequency rises, but it drops a lot more in some designs and materials than in others.

The highest power density and the highest speed don't emerge from the same materials and architectures, however. Between them, tantalum and ceramic capacitors provide the best performance, but only when they're used together. Solid tantalum electrolytic capacitors store the most charge per unit volume, but have higher ESR and are thus slower at the highest frequencies. They also tend to short circuit with over-voltage or high-pulse currents and tend to be very expensive. They can only be built so big, so banks of them are required to achieve higher power levels—which increases the overall risk of failure, and consumes more precious board space. Multi-layer ceramic (MLC) capacitors are generally the very fastest, because they have (at very high frequencies) the lowest ESR. But they too can only be built so big (maximum capacitance 10 to 100 times *less* than tantalums), and they are prone to cracking and catastrophic failure under thermal cycling. While faster and less expensive than tantalum, ceramics alone cannot begin to handle the power appetites of Pentium-class loads.

Delve into the innards of a ceramic capacitor, and you find architectures and techniques that are, in key respects, strikingly parallel to those employed in the manufacture of digital memory and logic devices. MLCs are built up of hundreds of layers of barium (and similar) titanates, each only tens of microns thick—architectural dimensions comparable to those commonly encountered in the semiconductor world. This design makes possible high capacitance per unit size, and surface mountable designs. One cost challenge used to be that ceramics were held hostage to the price of the palladium needed to form the electrodes. Palladium prices have surpassed the price of gold in recent years. The industry—Kemet included—quickly implemented suitable inexpensive substitutes, such as nickel or copper. Much of the future challenge in ceramics (and thus opportunity for new, clever designs) resides in making bigger units with higher capacitance (more and thinner layers) while remaining physically robust and reliable.

The somewhat slower tantalum capacitors meet that challenge for more capacitance by depending as well on equally fine-scale structures, but in this case resulting in a much greater robustness. Tantalums do so by employing a fine tantalum powder that's pressed into a pellet around a tantalum wire. The pellet is sintered to promote contact among the particles, creating a porous but electrically connected structure. Tantalum penta-oxide, formed on the exposed surfaces of the tantalum through electro-chemical treatment, is used as the insulator/dielectric. A cathode is formed in a dip-and-dry process, creating a thin film of manganese dioxide—lower ESR is obtained by employing an organic polymer cathode as an alternative. The resulting capacitance is very high because the dielectric layer is so thin, and the total surface area of the overall sponge-like structure of the anode is so large.

In the trade-off between raw speed, for which the MLCs excel, and power density, for which tantalum is best, the solution is to use both. Place the smallest, fastest power caches—the MLCs —closest to the final digital load, array the higher power density, but somewhat slower tantalums right behind them. The ultrafast ceramics provide a front-end tier of capacitance in units ranging from 0.1 to a few µF (with maximum reliable units pushing up to about 20 µF). Literally right behind are the somewhat slower (but by standards of all other capacitors, blazingly fast) tantalums providing the heavy lifting from 10s up to 1000 µF.

With both types of capacitors, there are delicate trade-offs to be made between using fewer, larger capacitors, and more numerous, smaller ones. Larger devices can store more charge—but they then require longer leads, which add inductance, which slows down delivery of charge to the final payload. Tens of these small capacitors are thus typically arrayed in banks around a Pentium or comparable digital load, to provide enough total capacitance while limiting the total inductance of the wires that link things together. Even as logic chips have grown bigger, more integrated, fewer in number, and more power hungry, the capacitors around them have grown smaller and more numerous. There is, now, something visually incongruous, about the huge microprocessor surrounded by comparatively tiny power hardware. From all outward appearances, the logic device looks large and cumbersome, now—the power devices surrounding it appear delicate and fast.

Kemet

Kemet was born in 1919 as a part of Union Carbide, to provide leading edge materials for vacuum tubes. The invention of the transistor in 1947 doomed that business—even as it spawned a new one for high-speed, high-power-density capacitors. By 1966, Kemet had emerged as the leading tantalum-capacitor vendor, and it remains so today. David Maguire, Kemet's current chairman and CEO, has been with the company for 43 years.

For most of those years, Kemet's edge was as much Wal-Mart as Motorola—and many market analysts still view the company in that light. Kemet maintains a huge 35,000 capacitor inventory. To serve its global customer base Kemet has embraced electronic interchange ordering and delivers worldwide within 72 hours or less with a remarkable 99+ percent on-time record. The company emphasizes quality control—capacitors are exposed to high electrical and thermal stress, and customers recognize that the capacitors are often the weakest link in an electronic system. Kemet's unglamorous skills in these areas are essential, but they do tend to obscure the high-tech race that's in progress under the hood of the capacitor itself.

Kemet produced its first ceramic capacitors in 1969. The company isn't anywhere close to the biggest manufacturer of ceramics today—but it's probably the best. Kemet's newest multilayer ceramics will hit mass production before the end of this year, with their flagship 475 layer unit (22 µF) built up from individual layers just 1.8 microns thick. Their next-generation device, already well along, is scheduled for release later next year and will have a record 750 layers (each 1 micron) and 47 µF. Up to 1000 layers may be possible before hitting the limits in manufacturing yield and physical reliability.

Much of Kemet's ceramics-related IP lies in continual refinement of the recipes and skills to manufacture these ever finer layers without losing yield or undermining device reliability. In addition, Kemet designs for itself much of the highly specialized equipment needed. Kemet's challenge is to do with layers what Intel does with gates—make them smaller and pack more of them into a single package. More gates provide more logic; more layers provide more power. Kemet is exploring technologies that will make large arrays physically stronger, even as the thickness of individual ceramic layers drops below 1 micron. This may require using semiconductor-fab equipment (e.g. vapor deposition), along with new types of dielectric materials.

Japanese ceramic capacitor companies, which had employed pure palladium electrodes, were the first to move away from that precious metal. Their U.S. counterparts, including Kemet, had long used a silver-palladium alloy and were therefore less seriously affected by the palladium price run-up noted earlier, and therefore reacted more slowly. Some of the trade press still makes much of the industry's migration away from palladium, but the companies themselves now view that as history. The problem with using other less costly metals like nickel and copper centered on the loss of capacitance during the process needed to bond these base metals to the ceramic, and thus the need to switch to more capital-intensive sintering equipment. By last year, Kemet had displaced roughly half of the palladium that would otherwise have gone into its ceramics.

Kemet's experience with tantalums extends back to 1958. Kemet works very closely with suppliers of tantalum to develop finer powders that increase surface area, and thus capacitance, lower resistance, and increase speed. Finer powder makes manufacturing more difficult, however—when the powders get too fine, the sintering stage tends to produce a formless blob instead of a fine, high-capacitance, microstructure. The mastery of such processes comes only from decades of experience—and Kemet has more relevant decades behind it than anyone else. Much of its know-how resides in trade secrets, but Kemet also has patents on the processes necessary to work effectively with ultra-fine tantalum powder, particularly those relevant to keeping impurity levels low, which is always a challenge with very fine powders.

Some of Kemet's cleverest innovations are now architectural. In March, Kemet introduced a record-breaking capacitor, the functional equivalent of a multi-layer tantalum, which has the world's lowest ESR for a tantalum device, six times better than the best that preceded it. This summer Kemet will release an improved design that cuts ESR another 30 percent, and which will thus approach the lightning-speed capabilities of the ESR of ceramics, while supplying 1000 µF of capacitance, at least 50 times more than the biggest practical ceramics. Kemet has reduced ESL substantially too, by using clever mechanical/electrical design to shorten the total wire paths without compromising capacitance or mechanical integrity—the company mounts several square capacitors side-by-side, in the same package, on a shared tantalum bus. Capacitors built to this design are relatively expensive, but provide a lower overall cost as they can replace multiple capacitors on a circuit board, saving board space and manufacturing costs.

The high-end of the capacitor business is part material, part manufacturing, and part circuit design

Kemet has also developed conductive organic polymer electrodes for its tantalum capacitors to replace manganese dioxide—the payoff is a much lower series resistance (again, the industry's lowest), and thus much higher speeds. Without the multi-layer (technically, a "multi-anode" structure) noted above, the single anode organic tantalum is slower, but it's cheaper. So in the hierarchy of capacitance needs, it fits another niche and increasingly displaces the less reliable large ceramics in many applications.

As the multi-layer tantalum illustrates, the high-end of the capacitor business is part material, part manufacturing, and now—increasingly—part circuit design. At the very high frequencies and power densities that the leading-edge capacitors must attain, the architecture and wiring becomes as challenging as anything encountered in the realm of high-frequency radio circuitry, for example. The capacitors must interact so closely and precisely with downstream loads—CPUs in particular—that capacitor designers now collaborate directly with their customers' design engineers.

Kemet has likewise been an industry leader in the shift from traditional capacitors (with wire leads) to surface-mount devices that are suitable for automated pick-and-place manufacturing, and that can be crammed much more densely on to circuit boards. The costs of expanding surface-mount capabilities comprise the major share of Kemet's $550 million in capital expenditures over the past five years, with more coming. Surface-mount capacitors now account for 90 percent of the company's sales. In 1998, Kemet formed a technical alliance with NEC to advance surface-mount technology. While Kemet had developed its own surface-mount device a decade earlier, the alliance with NEC brought in additional technical skills; NEC developed the world's first resin-molded surface-mount tantalum capacitors in 1981.

Kemet certainly squares off against a number of serious competitors. Among major domestic competitors are AVX (70 percent owned by Kyocera) and Vishay (which also manufactures precision resistors and powerchips). EPCOS (European), NEC, Hitachi, Murata, and TDK are others to contend with. Japanese firms do already dominate in the arena of low-power but high-speed capacitors for consumer products and the consumer end of telecom (cell phones, PDAs). But for the highest power density and speed combined, the U.S. players, with Kemet at the top, remain comfortably in the lead—we suspect in large part because many of their skills have emerged in meeting the requirements of the U.S. military and aerospace industries.

A second dimension of competitive concern for Kemet is that another company might gain a significant lead in manufacturing high-performance capacitors from some material other than tantalum. The spike in tantalum prices caused by the recent datacom bubble is history now, but while they lasted, sky-high prices spurred a vigorous search for substitute materials—particularly solid aluminum and niobium. Kemet itself has been studying niobium for forty years, and quickly brought to market last year a viable line of niobium capacitors. But niobium is much more electrically leaky than tantalum (4 to 10 times worse) and thus requires much thicker oxide layers. Kemet's niobiums work just fine, but the technology offers no compelling advantages over tantalum, which continues to gain market share.

Kemet has likewise made significant advances with solid aluminum technology—it began shipping its own line of fast aluminum capacitors in late 2001. The slightly slower, but much cheaper, high-power aluminums are finding rapid acceptance one tier down in the speed hierarchy of digital demand, and here too are increasingly displacing the large (fragile) ceramics; Kemet's best 100 µF aluminum with organic electrodes can hit an ESR of 10 milliohms, which makes it enormously fast for devices of this design, yet still very much cheaper and more robust than the ceramics against which it begins to compete. Kemet manufactures a multi-layer polymer capacitor as well, which employs metallized electrode construction. The electrode plates are composed of 100-300 Angstroms of base metal vapor deposited on a thin film layer of polymer; the plates are either wound, burrito-like, or stacked like pancakes.

Kemet has developed solid supply relationships with the most important electronics companies—computer manufacturers (e.g. Compaq, Dell), communications hardware manufactures (e.g. Alcatel, Nokia), suppliers of automotive electronics (e.g. Delphi, Bosch), providers of electronics management services (EMS) (e.g. Celestica, Flextronics), and the major distributors of electronic components (e.g. Arrow, Pioneer). As noted, Kemet formed a technical alliance with NEC in 1999 to produce high-capacitance, low ESR, organic polymer tantalum capacitors, and the company has a similar technical alliance with Showa Denko K.K. to produce solid conductive polymer aluminum surface-mount capacitors.

Logic and Power: Mirror Worlds

The diminutive logic gates have power appetites so minuscule one can almost count the individual electrons that flow back and forth to flip each individual gate or store each individual bit. Incongruous though it may seem, the operation of these triumphs of miniaturization depends on mirror-image arrays of powerchips (current gates) and power-caching capacitors, to assure the steady supplies of electrons that define data and conduct the logical calculations. The logic gates run at microamps and microwatts. Just millimeters away, the power circuits have to serve up currents that are eight orders of magnitude higher, but that change at comparable speeds. As we wrote last month, the power must—ultimately—be as fast as the logic elements, which means about ten times faster than the chip's clock speed; that's about how fast a gate has to flip so as to appear always "on" or "off," and never in between.

All flows of information are flows of power. "Seeing" is the art of sensing power as it flows from there to here. Morse's telegraph ran on batteries. The early telephone networks and switches depended on new electrical grids to transmit voice. One of the several "difference engines" conceived by Charles Babbage in the 1820s was going to be powered by steam. The ENIAC computer of 1946 consumed 174,000 Watts of electrical power to light its 18,000 vacuum tubes.

Might we eventually push the bit efficiency of our chips so high that power requirements become inconsequential? Physicists have pondered this question for several decades, but have yet to resolve anything of any practical interest. For a long time they believed that the minimum energy required to perform a single logical operation could not fall lower than the background level of thermal noise that exists in any physical structure (a quantity directly related to temperature T and expressed as "kT"). A landmark paper published in 1961 began to put this notion on a formal footing (R. Landauer, Irreversibility and Heat Generation in the Computing Process, *IBM J. Res. & Dev.* 5, 183—191 (1961)). It has since been established that—in theory—it is possible to build a computer that is not based on Boolean logic, and such a computer could—in theory—perform calculations without consuming any net amount of energy.

The faster the logic, the closer the power supply has to be located to it

But—an all important qualification—this theoretical result depends on the further assumption that time is of no object. Our hypothetical, power-free computer can run as slowly as it likes; it can take forever, if it must, to spit out its answer. To our knowledge, no one has yet developed a formal theoretical analysis for the minimum energy requirement to perform a logical operation within any fixed amount of time. A Nobel Prize probably awaits the young physicist or mathematician who eventually does.

Until that happens, practical engineering realities will establish what pure theory hasn't yet nailed down. To build a smarter microprocessor, you have to build a faster one, and to build a faster one you have to build smaller gates, and to run smaller gates you have to *lower* the voltage applied to each individual gate. To power the smarter, faster, higher-gate-count microprocessor, however, you have to *raise* the voltage somewhere back up the line, in the circuits that keep it lit—there's no other way to pump in enough current. So there's a zero-infinity schism looming here, between the discrete logic element and operation, which heads south toward zero voltage and zero power, and the aggregate power consumption of the intelligent chip, which heads north toward infinite voltage and infinite power. The faster the logic, the closer the power supply has to be located to it. The practical limits will ultimately be defined by material limits on just how far this close conjunction of high power and low logic can be pushed.

The biggest long-term challenge facing chip engineers—how to bridge the high-low chasm, how to push more power, faster, into larger aggregations of faster, lower-power gates. No one has yet worked out just what limits the laws of physics impose here, but they certainly do exist, and they are probably defined by the basic physical properties of materials. And for now, at least, that interface is largely defined by one device—the capacitor—two key materials—titanates and tantalum—and one corporate nameplate—Kemet's.

Peter Huber and Mark Mills
May 1, 2002

The Power Panel

We began our exploration of digital power technologies in September 1999, some 32 issues ago. The time has come to winnow down our Power Panel a bit, to emphasize the most fecund opportunities we've explored during that period—the areas where power technology is evolving the fastest, and having its greatest impact.

In brief, those technologies are the ones that lie on the right side of the great divide, between technologies dominated by the thermal and mechanical laws of the macrocosm, and technologies dominated by the quantum physics of semiconductor junctions. The technologies of greatest interest tend to be directly anchored in what the great physicist Richard Feynman characterized as "room at the bottom"—meticulously engineered atomic-scale junctions.

As we have argued from the outset, the stealth revolution in power is centered on the rising power levels that semiconductor technologies are able to handle. Silicon powerchips now switch and process kilowatts of power, not just milliwatts of logic or memory. Laser diodes and RF MMIC chips transform and amplify a thousand times more power than they did just a few years ago. Sensors detect all the various attributes of power with sensitivities 10,000 times greater than they used to. The leading digital power companies are taking advantage of thirty years of advances in semiconductor materials and chip-scale engineering, and trillion dollars of capital already invested to advance the manufacture of logic devices.

Quantum power technologies were our exclusive focus from September 1999 through June 2000. Then, before returning to our quantum interests in November 2000, for four months we ventured back into the world of Newton and Carnot, to address turbines, flywheels, and large fuel cells. We surveyed a number of excellent companies in those six issues: Beacon Power and Active Power (whose spinning hardware depends heavily on digital power electronics); Proton Energy, FuelCell Energy, Capstone Turbine and Catalytica Energy Systems (who likewise used digital power electrons to supply high-9s electrons, but whose businesses also remain highly dependent on the whims of regulators), ABB (a venerable Swiss technology leader, but one too prone to European daydreaming about green technologies), and GE, which still builds great turbines but which neglected to heed our (presumptive) advice when we told it to get rid of its captive bank. We still admire many of the companies and their key power technologies. But they just aren't at the center of the action that interests us the most. Accordingly, while we wish them all well, we no longer include them on the Power Panel.

Ascendant Technology	Company (Symbol)	Reference Date	Reference Price	5/01/02 Price	52wk Range	Market Cap
Electron Storage & Ride-Through	Kemet Corp. (KEM)	5/1/02	19.63	19.63	13.85 - 22.50	1.7b
	Wilson Greatbatch Technologies (GB)	3/04/02	25.36	27.50	23.00 - 39.00	574.0m
	C&D Technologies (CHP)	6/29/01	31.00	23.39	16.35 - 38.60	607.2m
	Maxwell Technologies (MXWL)	2/23/01	16.72	9.49	5.81 - 22.50	97.1m
	American Superconductor (AMSC)	9/30/99	15.38	7.50	6.50 - 27.90	153.5m
System Integrators	Amkor Technology (AMKR)	4/2/02	21.85	20.31	9.00 - 26.24	3.3b
	Emerson (EMR)	5/31/00	59.00	53.71	44.04 - 72.09	22.6b
	Power-One (PWER)	4/28/00	22.75	8.26	5.32 - 27.35	652.9m
Project, Sense, and Control	Danaher Corp. (DHR)	1/29/02	61.56	71.60	43.90 - 75.46	10.8b
	FLIR Systems (FLIR)	1/9/02	41.64	39.26	13.65 - 59.50	656.4m
	Analogic (ALOG)	11/30/01	36.88	48.03	33.40 - 56.50	634.0m
	TRW Inc. (TRW)	10/24/01	33.21	55.54	27.43 - 55.18	7.0b
	Raytheon Co. (RTN)	9/16/01*	24.85	42.40	23.95 - 43.05	16.9b
	Rockwell Automation (ROK)	8/29/01	16.22	21.51	11.78 - 47.20	4.0b
	Analog Devices (ADI)	7/27/01	47.00	36.80	29.00 - 53.30	13.4b
	Coherent (COHR)	5/31/01	35.50	31.47	25.05 - 40.60	902.5m
Powerchips	Cree Inc. (CREE)	4/30/01	21.53	12.49	10.59 - 36.65	906.9m
	Microsemi (MSCC)	3/30/01	14.00	13.19	12.48 - 40.10	376.0m
	Fairchild Semiconductor (FCS)	1/22/01	17.69	27.44	13.76 - 32.03	2.8b
	Infineon (IFX)	11/27/00	43.75	18.50	10.71 - 43.90	12.8b
	Advanced Power (APTI)	8/7/00	15.00	13.05	6.50 - 18.00	134.4m
	IXYS (SYXI)	3/31/00	6.78	8.59	4.27 - 19.45	230.5m
	International Rectifier (IRF)	3/31/00	38.13	46.97	24.05 - 69.50	3.0b

** 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.*

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

KEMET®

High Performance Solutions

Reprinted with permission by Gilder Publishing, LLC