THE HUBER MILLS digital power report

Special Report

The Powerchip Paradigm II

Broadband Power

by Peter Huber and Mark Mills

THE HUBER MILLS digital power report
Powering the Telecosm
SPECIAL REPORT
Published by Gilder Publishing, LLC

The Powerchip Paradigm II
Broadband Power

Bits will move electrons – the one and only form of broadband power – and move them so well they take over the click-click bang-bang infrastructure of the industrial age

"Eppur si muove" ("nevertheless, it does move") Galileo declared, after his conviction. He dared to assert that it is the earth that rotates, not the sun. The ecclesiastical court disagreed.

Let us now follow bravely in Galileo's footsteps. Gigabytes and terabits notwithstanding, atoms still move. Food and shelter, cars and planes, steel and silicon – most of our economic lives – still revolve around hard-ware, not soft. Oh yes, bits profoundly change and improve the way we handle the solid stuff of life; they even substitute for much of it at the edges. But at the end of the digital line there is still stuff. Most of our time, effort, and resources still go into moving atoms, not bits. We love the Web, but we still vacation in Tuscany.

And the movement of atoms is still ruled by the old laws, the laws of the macrocosm, Newton's laws, the laws of action and reaction, force and momentum, friction, dispersion, and decay.

Consider, for example, how we move our most cherished material assets – ourselves – from here to there. Like all other matter-moving technologies, the automobile must control position, speed, and acceleration, with sufficient precision to deliver the goods intact and on schedule. That takes layer upon layer of matter-moving technologies under the hood – pistons, shafts, disks, rotors, gears, links, belts, pulleys, throttles, valves, pumps, and fans – and array upon array of mechanical and hydraulic links, that push and pull, open and close, crank, pump, spin, reciprocate, and flow, all in precise synchrony.

It's the same in every plane, train, or truck, and on every factory assembly line. Webs of material-moving technologies are lined up, layered, and stacked. A General Motors (GM) assembly plant is a sprawling network of servers and pipes, architecturally quite similar to an Ethernet LAN, except that it takes a lot more than ether to roll a Buick off the end of the line. Same for McDonald's (MCD), Hyatt, Boeing (BA), United Airlines (UAL), or your local hospital, as it reshapes noses and removes gall bladders.

And same even – indeed, especially — in the chip fabs, the cyber-manufactories. No functioning PC emerges without the most fantastically precise control of the position, speed, and acceleration of silicon itself, and the systems that dope, etch, and package it. Stepper motors move wafers in ultra-fine increments. Chip masks are aligned to sub-micron levels of precision. Disk drives confront the strictly physical challenges of spinning platters, moving read/write heads, packing magnets on metal, and aligning laser pits on plastic. The end of Moore's law will be defined by the limits of our ability to process atoms, not bits.

However digital and virtual things may get in the middle, they will remain forever material at the beginning and end of the line. At the edges, it is atoms that move, not bits.

Power Train

It takes power to move them. Lots of it. Far more power than it takes to move and process bits. As we have emphasized from the beginning, the bit-moving business requires substantial amounts of very high quality power in its own right. But for all that, the moving of atoms requires considerably more. Roughly two-thirds of all the energy we consume is used to run motors – combustion motors, mainly for transportation, or electric motors, which do most of the heavy-lifting in less peripatetic applications.

So, getting down to the nitty gritty where the real action is, what *kind* of power does it take, to move atoms? Anywhere above the scale of molecular diffusion, there are only a few fundamental options. Power is conveyed by a mechanical force or torque that moves a translating or rotating shaft, or by pressure that moves a hydraulic or pneumatic fluid. Or by voltage that moves a current, that sets up a magnetic field, that creates a force on a shaft. There are many other ways to store energy, and storage systems can be carted around in various ways. But the essence of power is near instantaneous transmission, and that is where the most fundamental engineering challenges arise.

Electrons supply, by far, the highest energy density for power transmission. Compared to any mechanical or fluidic alternative, an electric wire offers effectively infinite power "bandwidth." A 740 kVA aluminum power line conveys about 1,000,000 kW of power through a cross section of a few square inches – about as much power as one finds under the wings of ten jumbo jets. (A superconducting power line boosts energy densities ten-fold higher still; a semiconductor runs an order of magnitude or so lower.) The steel power train in a Buick uses about five times the cross-sectional area of the power line to transmit a mere 100 kW to the wheels.

Electrical power is also, by far, the fastest and most responsive. Compared to all mechanical and chemical systems, electrical power flows can be turned on and off instantaneously. Electrical flows respond at close to the speed of light; mechanical systems are limited by the speed of sound.

Finally, it is much easier to add intelligence to an electrical power train than to any other. Electricity itself can be used to switch electricity – this is the essence of the transistor and all silicon smart chips and powerchips. In mechanical, hydraulic, and pneumatic systems, by contrast, the "switching" again occurs at Newtonian speeds. Logic comes down to "click-click bang-bang." Linkages, rocker arms, contacts, and valves flap back and forth, and are typically stopped by colliding with something else. Elaborately designed arrays of levers, gears, wheels, and cam shafts not only transmit the power but also impose upon it a desired trajectory and timing – just as the insides of a fine mechanical watch both power and impose a precise timing on the hands that move around the watch face.

The electrical power train is, in sum, overwhelmingly superior. It offers power bandwidth five orders of magnitude higher than mechanical or fluidic alternatives. Transmission speeds five orders of magnitude faster. Switching frequencies five orders of magnitude higher. It is broadband. And with suitable silicon control, it can be digital.

Power Transmission

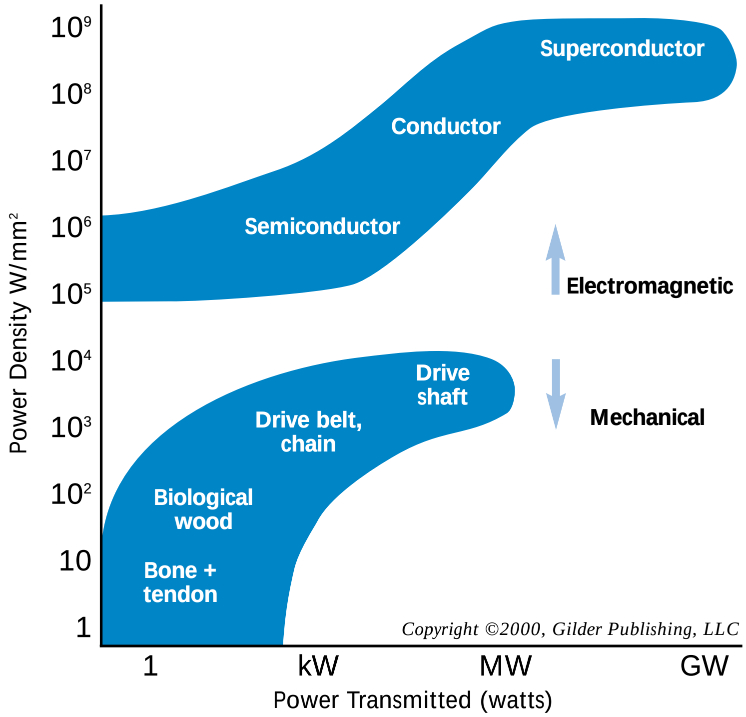

Power densities in mechanical power trains are ultimately limited by the physical strength of materials and the speed of sound. Power densities in electromagnetic power trains begin two orders of magnitude higher than the highest-power-density mechanical systems, and go up from there. They are ultimately limited by the dielectric strength of the conductor and the speed of light.

Dumbing Down

So why do mechanical and fluidic power trains still dominate so much of the power train universe? The answer, ironically enough, is that electric power is often too fast and too concentrated for control systems to handle. Or at least it was, until very recently.

How about all-electric steering in a car, for example? It's easy enough to make a compact "linear servo motor" that will push or pull a steering rod on command. Just send the right current, through the right wire winding, in the vicinity of a suitably mounted magnet, and the rod will move on command. That's basically how all electric motors work, and linear ones aren't very different from the more familiar rotating ones.

Except for their tendency to jitter or overshoot, to jerk out of control, or to fall off the edge – except, in short, for their horrendous dynamic instability. Unless they are carefully combined with inertial and frictional components to deliver something simple and steady – like rotating a shaft, which an ordinary electric motor can do so well – electric drives are febrile and unstable. They're fast all right, but their position, speed, and acceleration are almost impossible to control both accurately and simultaneously. Most drivers, however, insist on just that in the power steering of their car. Hair-trigger response is useful, it turns out, only in the very steadiest of hands. With anything less than perfect control, somewhat sticky and slow is often better than frictionless and hair-trigger quick.

And that's been the fate, until recently, of the electric power train. It has been extended a long way down the atom-moving line, but it has always given way to heavier, slower, more frictional systems toward the end. The power train of the nineteenth-century factory was a spinning shaft that spanned the length of the building; belts and chains delivered power to the individual work bay. Electrical systems have long since displaced most of that power backbone — the twentieth-century factory distributes power the length of the factory floor, to motors in every work bay, in each individual lathe, drill, milling machine, and so forth. But the shafts and belts and fluids are still there, still in control of the "last mile" – make that the last meter or two, typically – of the power train. Direct-drive electrical systems – systems in which the electrical power train acts directly on the final thing to be moved – have remained the exception, not the rule.

In the real world, in other words, we have deliberately loaded inertia and friction around the end of the electrical power train to tame its response. Most of the time, the electrical power train terminates in a dumb pump or motor, which then gives way to a fluidic or mechanical power train that interfaces with the ultimate payload. Down at the end, the engineers need the friction, the inertia, the sluggishness of the macrocosm. Broadband power is too quick to handle. So we dumb it down.

In a car, where there's a mechanical engine rather than an electric grid immediately at hand, it generally made no sense to go through an intervening electrical stage at all. The powertrain was mechanical at the front end, and had to be mechanical at the back end – so it was generally easiest to stay mechanical in the middle. Hence, the writhing snake's nest of shafts, belts, pulleys, chains, wheels, gears, and calipers under the hood. The Model T didn't even have a battery; the modern Buick remains a low-power 14 V electrical platform. The electric power is for start-up and ignition, and for distractions like power windows and the CD player. The real powertrain is entirely mechanical.

Old Powertrain

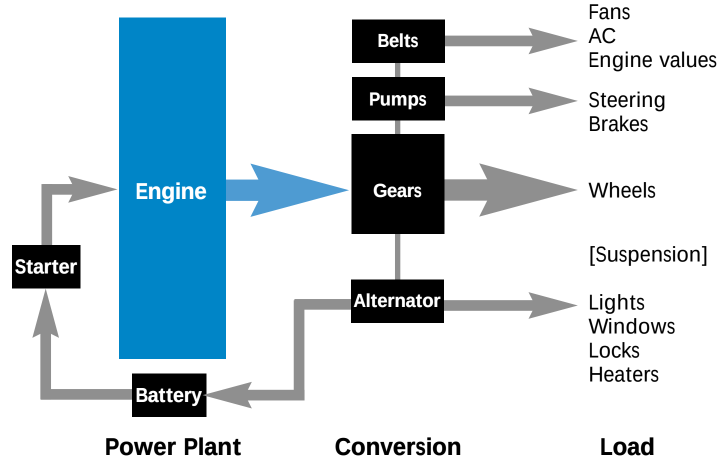

Four power transmission systems, two mechanical, one hydraulic, one electrical, interconnect through shafts, gears, and belts. The intelligence resides in mechanical switches, fluidic valves, and precisely machined metal.

Controlling Power at the Speed of Light

But now we have smart chips and powerchips – technology smart and fast enough to tame the hair-trigger power train, technology that can control light-speed power at the speed of light itself. This changes everything.

The electrical switch itself is the first and most fundamental piece of "motion control" technology to give way to the new silicon power technology. As we discussed in our first Paradigm issue, *The Powerchip Paradigm I: Digital Power*, (September 1999, reissue December 2000) the essence of the old electric switch is a pair of metal contacts that come together to close a circuit, or move apart, to open it. A click-click process, in other words. It is so electrically noisy that it all too often degrades power quality to computer-crashing levels.

It's slow, too, which can sharply limit performance of systems that require a whole lot of switching. Until the 1960s, the telephone switch was still a huge structure, circuits were set up by reconfiguring tapestry-like arrays of small, electromechanical switches – thousands and thousands of them, clicking away, day and night. A phone company end-office sounded like a warehouse filled with over-wound grandfather clocks. Nobody aspired to route terabytes of data through switches like those.

By the 1960's, however, the end of all that was already clearly in sight. Three Bell Labs scientists had won a Nobel prize just a few years earlier for inventing a new, solid-state substitute for the click-click switch –

the transistor. The rest, as they say, is history. Here and elsewhere, silicon switches have since pushed electrical switching speeds up into the gigahertz (at the micro-watt powers found inside microprocessors), or at least megahertz (for powers up into the kilowatts), or at least a kilohertz (for powers as high as megawatts). The phone or data traffic that it takes to order a pizza is routed end-to-end through silicon switches, not electromechanical ones. The rotary dial telephone – a click-click signaling structure – is gone too. Click-click is finished.

Until the pizza delivery man hops into his car that is – there, the click-click bang-bang still rules. "The rest," in other words, isn't all history by any means. Not yet. A huge compendium of silicon history remains to be written, in all the atom-moving sectors of the economy where most of the power is consumed.

It isn't particularly difficult to build a direct-drive servo motor, and that half of the technology has been around for a long time. The "servo" implies, quite simply, a feedback loop between the motor and its power supply. A controller keeps track of what's actually happening at the far end of the motor, and continuously updates the power feed to conform that to what's supposed to be happening. The "direct-drive" part just means that the motor is right on the doorstep of the final payload, rather than separated from it by some elaborate array of intervening click-click bang-bang mechanical or pneumatic linkages. The servo motor can be rotary, like a conventional electric motor, or it can be linear, like a conventional solenoid. Either way, it can execute choreographed moves over a specified trajectory, speed, and acceleration. Just feed it power of exactly the right profile, and it will move where you want it to – much as a loudspeaker will vibrate its way through a Beethoven symphony if you send it the right electrical signal. To move something in three dimensions, or to combine linear motion with rotation, use multiple servos, in parallel or in series.

How is this fundamentally different from how electric motors are commonly used today? The electric power is now converted into mechanical force only at the very threshold of the payload. Click-click bang-bang mechanical control systems give way to digital electric ones. And the intelligence now resides upstream of the motor, not downstream. The power is precisely contoured and shaped before it is converted into force and motion, rather than after.

And this is where the key breakthroughs have recently occurred: in the contouring and the shaping, in the units that dispatch just the right amount of electric power, instant by instant, to the servo motors. Until recently, direct-drive motors had two upstream units, a "controller," that did the thinking, and an "amplifier" that boosted the controller's signal to the power level required to drive the motor and its load. Today, the two have collapsed into a single box – a silicon driver packed with smart chips (for intelligence) and powerchips (for power). Smart chips in the drive work out, in real time and at very high speed, what kind of power profile to send to the motor's coils to move the payload along any defined trajectory. Powerchips build the profile itself. The motor constantly updates the drive on how it's doing. And the power profile is continuously updated. The direct-drive motor has been tamed, in short, by the confluence of smart chip and powerchip in a single high-intelligence, high-power digital box.

The servo motor can be rotary or it can be linear. Either way, it can execute choreographed moves over a specified trajectory, speed and acceleration.

The servo motor's dynamic stability is maintained by a very tight ("high gain") feedback, and a tremendous amount of real time computation. Stability depends entirely on extremely fast response to the feedback. The drive now lacks the gears, ballscrews, fluid reservoirs, and other inertial and frictional devices that indirect-drive systems rely on to damp out disturbances and unexpected variations in the load, like a rough spot on tarmac under the wheels, for example. "Induction" servo motors – in which the motor's own magnet is itself created ("induced") on the fly, by current flow through the system — require the most control of all, because here the configuration of the motor itself keeps changing along with everything else. It's like trying to balance a straw, or maybe a piece of spaghetti, on the tip of a finger – if you can do it at all, it's only by watching closely, concentrating hard, responding very fast, and thinking ahead. Accurate, smooth control of position, speed, and acceleration often requires update intervals in the milliseconds for position, and in the microsecond for velocity and acceleration.

The control algorithms are so demanding, that until recently they required far more computing power than could practicably be supplied, at least near the places where the old click-click bang-bang hardware held sway. GM didn't have room for a mainframe under the hood of the Buick, and computing power cost too much in any event. But in the last few years, smart chips have emerged that are easily up to the task. Motion control software now runs on Pentium-class microprocessors – though driver manufacturers generally use their own custom-designed integrated circuits. Increasingly, motion control smart chip and powerchip are physically co-located on the same silicon substrate, or at least glued tightly to each other. The new digital drives run very fast with powerchips operating from the kHz to the MHz range depending on the applications.

Equally essential have been the dramatic improvements in powerchip technologies. The GHz smart chips can calculate exactly how to shape the power delivered to the servo motor downstream, but they can't create the power profile itself, not even for disk-drive, still less for a one kW motor-driven factory-floor robot, the warehouse fork lift with its 20 kW servo motor, the 2 MW motors propelling a locomotive, or 5 MW motors pressurizing interstate natural gas pipelines. That takes powerchips that can handle muscular currents and voltages, and handle them fast enough to synthesize exactly the right power transient. As discussed in our April 2000 DPR, and further in the December 2000 issue, the powerchips have recently come of age, too. Performance and reliability have improved dramatically and costs have plummeted, from around \$0.50/W in 1990 toward \$0.10/W now in sight.

Standardized, modularized units now provide a common drive hardware and software platform for all motor types – rotary or linear, induction or permanent magnet. The enormous diversity in application, power levels (covering seven orders of magnitude), speed, and performance in the atom-moving world has lead to at least a dozen major communication and interconnection "standard" networks, all with significant market penetration, including LonWorks, ControlNet, Fieldbus, Seriplex, ARCNet, AS Interface, Interbus, Profibus, and SERCOs. The standard networks create the capability (within broad applications areas) for cross-vendor and function uses and interconnections of digital controls, drives, actuators and CPUs.

Some systems are hard-wired controls for specific applications; others are fully software controllable. Industrial robots – which are, by and large, just elaborate arrays of electric servo motors – can now be instantly reconfigured to perform new tasks through software alone which is a dramatic advance over previous systems that often required hours of manual adjustments.

The very rapid advance of direct-drives has propelled parallel, if somewhat less dramatic, improvements in the servo motors themselves. Prices have dropped steadily. Magnetic materials and bearings have improved. Brushless, permanent-magnet synchronous designs still dominate; linear-induction asynchronous motors are on the rise; switched-reluctance and voice-coil designs are used too. Different designs have been optimized for positional "stepping" or more complex choreography, stroke-length, speed, acceleration, maximum force, operating temperature, compact size, durability, electrical efficiency, manufacturing simplicity, and cost.

And there have been parallel advances in sensor technology, which provides the essential feedback from the servo motor to the controller. Fiber optic and even wireless buses are now routinely used to provide very high speed and robust communication between sensors and controllers, and among the arrays of controllers it takes to run an entire assembly line, an electric locomotive, or control a jet.

What it all comes back to, in the end, is the technology of high-9s power. In our original *Powerchip Paradigm I: Digital Power*, we discussed what kind of power it takes to move bits – and the answer is perfectly smooth sinusoid (AC) or flat-line (DC) power, reliable down to the nanosecond. Now we have the *Powerchip Paradigm II: Broadband Power* – the kind of power it takes to move atoms, not the old way, via narrowband mechanical and fluidic systems, but the new way, by way of direct, broadband electrical drives. The high-9s power doesn't have to be quite as fast here – millisecond-level accuracy will generally suffice. But in other respects it has to be massaged even more delicately – the objective now isn't a sinusoidal AC or a flat-line DC, but rather is a continuously varying power curve, redefined instant by instant to achieve a specified force, torque, position, speed, or acceleration at the far end of the servo motor downstream.

New Powertrain

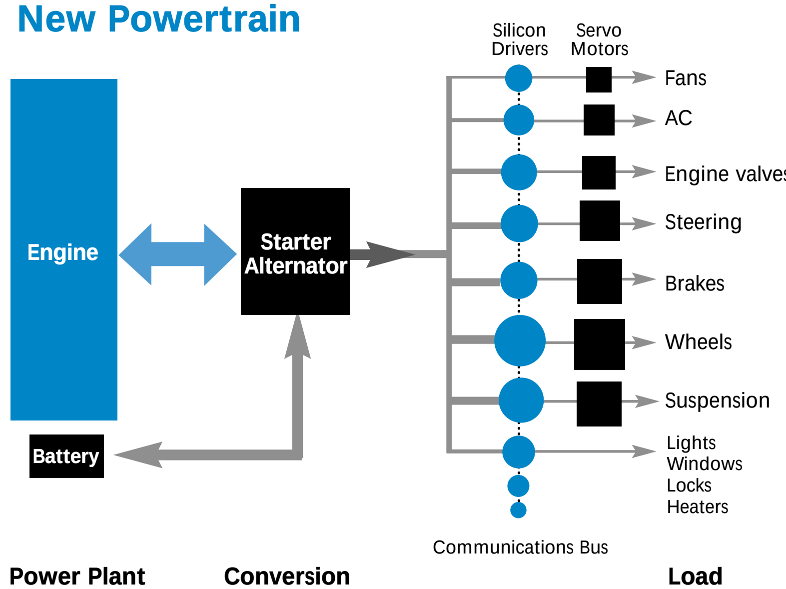

An integrated electrical power transmission system, terminating in silicon drivers and servo motors, interconnects by way of a digital communications bus. The intelligence is lodged in silicon smart chips and powerchips.

High-9s Motion Control

The pay-off, if you can do that just right, is digitally programmed motion. Newton's laws reincarnated in the digital, broadband power train. The silicon computer upgraded to general purpose silicon drive — a digital mover not only of bits, but of atoms too.

The first and most immediate pay-off is motion control of precision unattainable by any other means. Linear motors can provide position control of submicron precision — an order of magnitude better than conventional solutions.

This is why they have already completely taken over applications that require ultra-fast and/or ultra-precise control – the manufacturing of chips. Many of the early and best applications of linear technology are in semiconductor and electronic processes – for wire bonding, wafer probing, laser trimming, pick and place, and printed circuit board drilling and inspection. Servo motors are used exclusively to spin disk drives and to move the read/write heads.

Slicing up a silicon wafer requires a motor that will move the crystal in very precise increments. The conventional technology for moving an object along a linear track is a rotary-driven actuator. But it's performance is limited by friction, wear, jitter, backlash, hysteresis, temperature effects, and so forth – all of which can readily be compensated for using a direct-drive servo motor and dynamic feedback. The actuator itself no longer establishes the outer limits on positioning precision. Direct digital drives have likewise taken charge of position control in medical analysis and treatment (X-rays and MRI) and precision machining.

Response speed is the second big pay-off. Because they move less material in the middle, direct-drives have far less inertia and friction; and because they are controlled by computers, they can react much faster to the outside world. To put it in engineering terms, they have much higher natural frequencies – they are stiffer. This used to be part of the problem – this was what made direct-drives febrile and unstable. But with digital controllers in charge, high stiffness is now better, not worse. Supersonic jet fighters with forward-swept wings are far more maneuverable than conventional designs precisely because they are dynamically unstable – so unstable that they remained wholly impractical until digital power systems took over the flight controls.

Direct-drive digital systems are smaller and lighter, too, because electrical power requires so much less transmission hardware than mechanical or fluidic alternatives. This is an advantage even on a stationary assembly line, and a very substantial one when it shaves hundreds of pounds off the weight of a car or plane. The electric calculator contains much less stuff than an abacus; an ink-jet printer far less than the click-click bang-bang daisy-wheel printer you owned fifteen years ago, or the IBM Selectric a few years before that; and the Pentium on your desk far less than the all-mechanical computer conceived by Charles Babbage in 1822. A similar, and equally dramatic slimming down is now under way in mechanical systems across the landscape, including much larger ones. Including those under the hood of the car.

Silicon Driver

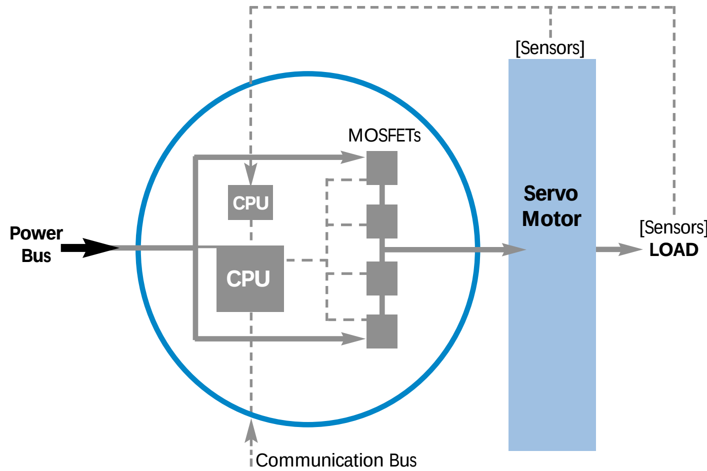

Redundant silicon microprocessors plus an array of MOSFET powerchips modulate electric power to control the position, speed, and acceleration of the payload. A high-gain feedback loop links sensors at the load to the microprocessors; a communications bus connects multiple silicon drivers on a car-wide web.

Finally, with less stuff in them, and far fewer moving parts, direct-drives are generally more robust, and less environmentally sensitive, too. Pneumatic and hydraulic fluids leak, turn into molasses when they get cold, and are easily contaminated. Mechanical links need lubricants and can get bent out of shape when they expand or contract. As a power-transmitting technology, an electrical wire can be far more robust, and far more tolerant of hostile environments, and much less in need of periodic maintenance, than a shaft, belt, pulley, or fluid-filled pipe.

The future of electric drives everywhere else has already played itself out in the humble wristwatch. The mechanical chronometer of yesteryear was a thing of great engineering beauty, magnificently designed, meticulously assembled, cherished by its owner, and handed down as an heirloom from generation to generation. It was expensive, too. But today, you can find far better timekeepers in children's toys. The piezo-electric effect using a quartz crystal (cousin to silicon) has displaced the grandfather clock's pendulum and the wristwatch's mechanical hairspring, balance and lever system. Speed-of-sound accuracy –seconds per day – has given way to speed-of-light levels of accuracy – seconds (or better) per decade.

The Hollow Network

Broadband power is the technology that moves more bits and more electrons – but far fewer unnecessary atoms. It moves more final payload and less click-click bang-bang deadweight in between. The gains come from pushing electrons, instead of shafts or fluids, right down to the edge of the payload itself.

The effect of limitless bandwidth in glass is (ultimately) to hollow out the bit network, to push the digital intelligence out to the edges. The advent of limitless power bandwidth has a parallel impact on power trains. The painstakingly complex and "smart" arrays of click-click bang-bang belts, pulleys, gears, valves, rocker arms, and pulsing fluids – all disappear. Pure power – electrical power – is generated at one end of the power train, and is converted back to pure motion at the other end.

The systems in between are all electrical; no mechanical intelligence remains in the middle, because nothing mechanical remains there at all. Throw out all intermediate layers of mechanical control of matter; replace them with the far more compact and intelligent controllers of electrons. The matter in the middle is the dumb, electrical wire; its only function is to convey a lot of power through a very small conduit. The thinking is done in silicon – smart chips plus powerchips – which is almost equally frugal with space and weight. All the intelligence is directed at choreographing flow of electrons – down to highly responsive electromechanical interfaces – the servo motors – situated at the very end of the line, where the actual payload is located.

James Clerk Maxwell, father of the physics of electromagnetism, could not have dreamed of today's silicon-rich servo motor technologies a century and half ago, but he could readily have grasped their transformative potential. Nineteenth-century physics understood the core fundamentals that govern here, most specifically, the vast gulf between the speed of sound and the speed of light. Photons – light, electromagnetic waves, and electricity – travel five orders of magnitude faster than sound, and thus commensurately faster than any transmitter of force, torque, or fluid pressure. Similar ratios reappear in the power densities of electric wire versus steel shaft. And in the clock speeds (or stiffness or natural frequency) of mechanical- versus electrical-drive systems. Everything else follows from those fundamentals. Get the electric power under sufficiently tight control, and it will fantastically outperform any mechanical alternative.

With the silicon technologies now at hand, the tight control is finally there for the taking. Which means that broadband power will, step by inevitable step, overtake all the conventional, narrowband alternatives.

That's a big change – bigger than any other we've written about in his newsletter so far. It takes far more power to move atoms than bits. For every powerchip upstream, conditioning power from a UPS, flywheel, microturbine, or substation of the grid, there will be 10 to 100 more powerchips in the atom-moving macrocosm downstream. This is where the other 90 percent of the Powercosm will unfold – the Powercosm downstream of the smart chip, where bits orchestrate the movement of atoms, rather than upstream, where the atoms impel the movement of bits.

If the other 90 percent of the Powercosm will unfold somewhat more slowly, it's for reasons rooted in industrial history. The market for high-9s power crystallized first around computer and communications technologies, because these technologies can't run at all without it – when the electron supplies fail, the bits stop moving. Mechanical systems, by contrast, can move just fine with virtually no electrical power at all, and have been doing so since the days of Rudolf Diesel and Henry Ford. Powercosm technologies are taking over the mechanical world not because the wheels won't roll without them, but because the wheels can roll much better with them. The world of mechanical engines is at least a century old. Fundamental restructurings take longer in mature industries, especially when they remain optional in the short term.

But then, the same could have been said, and often still is, about optional microprocessors on the desktop, or optional bandwidth in the Telecosm. Who really needs it? The old Selectric works fine, and the telephone does too. Until the competition goes digital. Then the Selectric retires, along with its owner.

As we argued in September 1999, in *Powerchip Paradigm I: Digital Power*, the digital world is fueled by very stable, reliable electrons. The electrons move the bits. What we have argued here is that the bits, in turn, will move electrons – the one and only form of broadband power – and move them so well that they will take over the tasks formerly performed by the whole, sprawling, click-click bang-bang infrastructure of the industrial age.

The companies in the thick of this transition are going to get stupendously rich. *Eppur si muove.*

Peter Huber & Mark Mills

Editors	Peter W. Huber Mark P. Mills
Publisher	Lauryn Franzoni
Managing Editor	Debi Kennedy
Designer	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. Editorial inquiries can be sent to: peterhuber@gildertech.com or markmills@gildertech.com. Single-issue price: $50. For subscription information, call 800.261.5309, e-mail us at help@gildertech.com, or visit our website at www.Powercosm.com