GILDER
TECHNOLOGY
REPORT

November 2004 / Vol.IX No.11

Telecosm and the Real Promise of the Nanocosm

Even as AT&T, MCI, and Level 3 wilt amid the chronic American telecom slump, the Corvis-equipped Broadwing network is beginning to show its superiority

Inside:

As snow swirled through the Resort at Squaw Creek and around the new **Verizon** (VZ) base station in the parking lot, frosted the panoramic windows beside the Wi-Fi antennas from Tropos Networks, and whited out the local airport at Truckee, connectivity was king, as an issue and a promise.

It was the 8th Annual Gilder/Forbes Telecosm conference. A previous one was scheduled for September 12, 2001, and another came in the midst of the crash, but this was the first one held in a blizzard. Nonetheless, sponsor Tropos provided robust connectivity throughout the lobby next to the conference ballroom. Thus, after journeying from the Berkshires to the Sierras, our Charlie Burger journeyed once more, lugging his geriatric laptop from the ninth floor of the resort's tower wing—through lobby, snow and ice between buildings, past resort reception desk and atrium—to the telecosmic future on the far side.

Perhaps Charlie should have contemplated the fate of that party of pioneers who perished two centuries before in an early blizzard in nearby Donner Pass. Among them was one Charles Burger. Was he a distant relative? Who knows? But perhaps no Charles Burger is destined to succeed in California's Sierras. After discovering the Tropos network at last, his HP announced the absence of the Internet beyond. Through the fractured Windows, Charlie felt a chill gale of creative destruction.

Why hadn't Charlie heeded **Qualcomm** (QCOM) founder Klein Gilhousen's sage warning uttered prior to his final journey to failed connectivity? Who in his right mind, Klein asked, in an age in which the supreme scarcity is time, will run from coffee shop to laundromat to motel lobby to unpack in search of the Internet?

If Wi-Fi/WiMAX is to compete with CDMA in mobile connectivity, it must be everywhere. Otherwise, it might as well be nowhere. The savvy Tropos laptop meds recounted their successes in deploying antennas in parking meters and their partnership with government to "increase competition" in broadband access. As Corpus Christi, aided by Tropos, keeps private enterprise at bay by providing 20 square miles of "free" tax-supported connectivity, followed by larger cities such as Philadelphia and much of Utah with its $350 million Utopia broadband project, we see that WiMAX and Tropos may have found their Daddy Warbucks.

Meanwhile, according to Andrew Odlyzko, U.S. Internet traffic is growing 70 percent per year while traffic in the rest of the world is doubling yearly. Americans download 1,000 megabytes per person monthly while Koreans

CONTINUED ON PAGE 2

download 2,700 MB/person and Hong Kong 4,529 MB/person, also monthly. At 70 percent per year growth, it takes us almost three years to catch up to the Hong Kong of today. Meanwhile, with a yearly doubling of traffic, Hong Kong in three years might well be downloading 36,000 MB/person or over seven times our downloads at that time. Our key Asian competitors are not only ahead of us but may be pulling away in the near future, as we look to the regulators of roads and street lights to pull us through.

But none of this matters, so says the dashing Roxane Googin, editor of the *High Tech Observer*, also speaking at Telecosm. Author of "The Paradox of the Perfect Network," she has discovered a felicitous telecosmic business plan: charge investor clients $10,000 for newsletters and associated sources of advice to flee all investments in telecom. Googin tried this plan before, during the deflationary death spiral of 2000–02. It worked "perfectly" then, so why won't it work now? Perhaps because the election is over, hyper-regulation and price controls lost, and the Asians are showing us how to roll out broadband.

As the faithful looked through the falling flakes, what they couldn't find was the perfect network of commoditized bandwidth. What they could see was that as the marginal cost of transistors has trended relentlessly toward zero for decades, **Intel** (INTC) has continued to make money. They could also see that networks still require updating and maintenance, that broadband connectivity is in dearth—and that Francine Berman of the San Diego Supercomputing Center is still searching for a network that can transport her real-time, potentially life-saving information on brain tumors for surgical diagnostics. Dr. Berman, meet Dr. David Huber of **Broadwing** (BWNG), if you can reach him before he flees through the snows to make a meeting with a potential customer in Chicago.

Broadwing's bit blizzard

Even as **AT&T** (T), **MCI** (MCIP), and **Level 3** (LVLT) wilt amid the chronic American telecom slump, Huber's Corvis-equipped Broadwing network is beginning to show its superiority. Facing the fact that transporting bits the last 20 miles costs five time as much as transporting them over the backbone network, Huber acquired Focal Communications in September to complete the transformation to vertical integration required for success in an undershoot technology. Yes, telecom remains in undershoot, incapable of delivering the broadband services of a transparent global economy moving onto the Net. While the rest of the industry prattles about a fiber glut, Huber knows that the basic problem is a U.S. connectivity dearth and broadband paralysis.

In mid-November, Huber revealed that while his rivals continued to shrivel, Broadwing's sales held steady (before adding in Focal revenues) and data revenues perked up 5 percent sequentially. Focal added $20.5 million to the quarter, bringing total sales to $163.4 million. Domestic average termination costs per minute plummeted 54 percent year over year and lowered average cost per customer access circuit by 17 percent.

Broadwing's performance is a tribute to the technological edge imparted by Corvis, which has enhanced the cost effectiveness of optics (measured in wavelengths times bit miles per cable without optoelectronic regeneration) by 16 thousand-fold in eight years. Now capable of lighting 320 wavelengths per fiber, Broadwing is the only network with widespread deployment of Raman amplifiers, increasing the distance between electronic regeneration in the network from 400 kilometers to 2,000 km and enabling all-optical wavelength switching, a much cheaper and more reliable architecture than a packetized IP-routed **Cisco** (CSCO) core.

Enter **Xan3D**. In an historic presentation at Telecosm that daunted all the many doubters of his awesome claims, John Trezza of chip innovator Xanoptix (now Xan3D) defined a radical new path to ubiquitous broadband. Xan3D, he believes, will propel the WDM all-optical revolution with its fused-chip technology alternative to the system-on-a-chip (SoC) approach pursued by the rest of the industry. Not only will Trezza push the network edge closer to the end user, he will also open the floodgates of bits bottlenecked between chips and let them flow freely. Today, the interfaces between electronic storage and optical bandwidth remain beset with needless complexity, a jungle of conflicting systems and protocols linked by miles of microscopic metal wires on every silicon integrated circuit, feeding buses and converters, buffers and power drivers down the pins of chips or surface mounted devices and across printed circuit boards off to network interface cards on hubs and switches and routers leading to optoelectronic connectors and transceivers coupled to hybrid networks of optics and electronics. All this takes time and power.

Xan3D tames this jungle of complexity by laminating or implanting one chip upon another using a proprietary fusion process. This technique overcomes all the incompatibilities among different semiconductors and metals and links the chips down insulated vias on the sides that bypass the usual pin or ball connectors on circuit boards and backplanes. Thus Xan3D replaces all the power-hungry buses, pads and drivers, on-chip memories and off-chip transceivers with an integrated stack of chips. On top is an optical communications plane with up to 8 million contacts per square centimeter, enabling off-chip connections as fast as the on-chip bandwidth allows. The optical contacts can be vertical-cavity surface emitting lasers (VCSELs), edge-emitting lasers, or light-emitting diodes combined with an equal number of photoreceptors. Since optics still performs all the remote communications, the all-optical network prevails. All-optical networks will always be superior to optoelectronic networks in both latency and reliability, regardless of improvements in electronic hardware.

For Trezza, the future of optics is to take a six-foot rack of equipment with 100 - 200 modules and reduce the whole thing to the equivalent of a laptop in terms of complexity and performance. Even some passive all-optical processing will disappear from the network as advances in transceivers enable optical signals to travel farther down the fiber before degrading. There's no way to get any better or cheaper than by having just fiber between endpoints, Trezza reminds us. The entrenched MCI/ATT culture may be fooled into strewing Trezza's "laptops" all over the network, using the new optoelectronics to obviate the promise of optics. The possibility of a patchwork landscape of rack-mounted laptops may prove a tempting shortcut to the incumbents as they shun the all-optical dream and labor to extend the life of their existing infrastructures.

At every node, however, these bit-routed networks will still have to dig down through the digital hierarchy to unearth the individual packets. Xan3D hardware will run cheaper, simpler, and cooler per packet, but the number of packets in the network will multiply in proportion to the rate at which Trezza frees them from their electronic prisons. Thus, Trezza's advances in electronics will be matched by the growth in the flow of bits, and optoelectronic networks will continue to drown in complexity.

As Trezza pushes the optoelectronic edge closer to the end user and opens the floodgates of bits bottlenecked between chips and lets them flow freely, it will be Dr. Huber's challenge to channel this flood onto the Broadwing network. With his all-optical infrastructure already in place, Huber can add new wavelength circuits at virtually no cost. He has no alternative for survival; to realize run-away all-optical leverage over his electronic network rivals, Huber must light lots of lambdas. Thus, his real adversaries are not the AT&Ts of the world but the academic road hogs, government regulators, and promoters of perfect networks.

—George Gilder and Charles Burger

* * *

Reading the Small Print

Are you looking for nanotechnology? Well, nanotech refers to things you cannot see—anything for which important dimensions are measured in nanometers (billionths of a meter). The diameter of a silicon atom in a microchip or optical line is about 0.23 nanometers. By that standard, this is certainly a nanotech, or even a picotech (trillionth of a meter) newsletter. The storage domains on dynamic random access memories are measured in femtofarads of capacitance (that's 10^{-15}). Are DRAMs femtotech? Are the dimensions of a device what make it distinctive?

I have a photo of a sign that reminds me of nanotechnology. Huge letters warn: "CAUTION. This sign has SHARP EDGES. Do not touch the edges of this sign." Small print at the bottom reads: "Also, the bridge is out ahead." Nanotechnology is like the "sharp edges" part of the sign. Everyone is so focused on the big letters that the important part of the message goes unnoticed. The semiconductor industry is like the bridge warning. Let's return for a moment to the state of the semiconductor industry.

By 1991, transistors were already smaller than bacteria. By the 370-nanometer generation, transistors were smaller than the wavelength of light. Transistors cannot be seen with an optical microscope. This year, companies are building chips with 90-nm transistors. That's just twice the size of a *virus*. At 90 nm, 100,000 transistors fit on a small grain of sand. That's why big chips have hundreds of millions or even billions of transistors.

We like Cepheid, a biotech chip company that uses DNA to identify biothreats, and Microvision, a MEMS heads up display innovator

No surprise that the world of semiconductors is converging with the world of biological systems and with the world of atoms and molecules.

The gate oxide, the insulator between the transistor and its controlling gate, is the transistor's smallest dimension. In leading-edge semiconductor processes, this gate oxide thickness—perhaps we should call it thinness—is a mere three atomic layers, measurable in the hundreds of picometers. Some nanotech enthusiasts say this is not real nanotech because the structures are not built by molecular assembly. Semiconductors, they argue, merely whittle material to small dimensions while nanotech assembles devices at the molecular level.

Future semiconductors, however, will combine top-down etching with bottom-up assembly. And, it is the semiconductor industry that is the enabler for progress in molecular assembly. Semiconductor industry investment paid for advances in materials purification, in lithography, in materials science, in precision process control, and in advanced microscope design. The semiconductor industry paid for the development of the tools, equipment, and materials that support nanotechnology experimentation. And, semiconductor laboratories at **NEC** (NIPNY) and **IBM** (IBM) developed the most notable nanotech innovation, the nanotube, the central technology in a number of nanotechnology startups.

Semiconductor industry goes tubing

To understand what a carbon nanotube is, think of chicken-wire fencing. Chicken wire is a mesh of connected points forming hexagonal openings. Imagine that the

TELECOSM TECHNOLOGIES

Advanced Micro Devices	(AMD)
Agilent	(A)
Altera	(ALTR)
Analog Devices	(ADI)
Broadcom	(BRCM)
Broadwing	(BWNG)
Cepheid	(CPHD)
Chartered Semiconductor	(CHRT)
Equinix	(EQIX)
Essex	(KEYW)
EZchip	(LNOP)
Flextronics	(FLEX)
Intel	(INTC)
JDS Uniphase	(JDSU)
Legend Group Limited	(LGHLY.PK)
McDATA	(MCDTA)
Microvision	(MVIS)
National Semiconductor	(NSM)
Power-One	(PWER)
Qualcomm	(QCOM)
Samsung	(SSNLF/SSNHY)
Semiconductor Manufacturing International	(SMI)
Sprint	(FON)
Synaptics	(SYNA)
Taiwan Semiconductor	(TSM)
Terayon	(TERN)
Texas Instruments	(TXN)
VIA Technologies	(2388.TW)
Wind River Systems	(WIND)
Xilinx	(XLNX)
Zoran	(ZRAN)

Note: The Telecosm Technologies list featured in the Gilder Technology Report is not a model portfolio. It is a list of technologies that lead in their respective application. Companies appear on this list based on technical leadership, without consideration of current share price or investment timing. The presence of a company on the list is not a recommendation to buy shares at the current price. George Gilder and Gilder Technology Report staff may hold positions in some or all of the stocks listed.

ADVANCED MICRO DEVICES (AMD)

PARADIGM PLAY: INTERNET COMPATIBLE PROCESSORS

NOVEMBER 19: 20.81, 52-WEEK RANGE: 10.76 - 22.49, MARKET CAP: 7.64B

AMD wants to flood the world with cheap personal computers. On 28 October, AMD announced the Personal Internet Communicator (PIC). PIC, a Windows computer based on the Geode GX500, has a 10-GB hard disk, 128 MB of memory, and a modem. A complete system with a display is $249. This computer is for emerging markets outside the United States, where it will put AMD in direct competition with VIA. If building for volume is the right thing to do (Tredennick's law says it is), then AMD is doing just that by building for the low-end, rest-of-world (ROW) market. AMD's strategy at the high end is excellent (64-bit extensions to x86, dual-core Opteron, HyperTransport), its desktop strategy is excellent (dual-core Athlon, HyperTransport), and its low-end strategy is excellent (PIC). PIC is the first step in integrating Geode into AMD's strategy; next will be embedded applications. Embedded applications are already 60% of VIA's unit sales; AMD will be there too. Embedded x86 applications are the next frontier and they dwarf all other markets in unit volume. When the x86 invades embedded systems to take over those applications and displace ARM, MIPS, PowerPC, and hundreds of older instruction sets, AMD and VIA will be there to profit from it.

Broadwing (BWNG)

PARADIGM PLAY: THE PARAMOUNT ALL-OPTICAL COMPANY

NOVEMBER 19: 5.51, 52-WEEK RANGE: 5.11 - 30.70, MARKET CAP: 327.72M

Broadwing's 3Q results, which included one month of newly acquired Focal Communications, showed early signs that management's strategy both to lower network access charges and to take market share, may be bearing fruit. Total revenues of $163.4m included an expected $3.8m sequential decline in voice sales, offset by a $4.9m increase in data revenues; Focal contributed $20.5m to the quarter. Gross margin of 31%, up from the 20s last year, is expected to climb to 35% - 40% next year. Focal's local networks should add yet more to access savings while increasing top-line growth through cross-selling to Broadwing customers.

The balance sheet took a hit as net current cash fell $150m sequentially to $120m, due mostly to one-time charges related to the acquisition. This greatly increases the likelihood that Broadwing will continue to pay its $225m convert in stock. The company has repeatedly emphasized that a strong balance sheet is a major marketing tool with which to lure large customers. In a worst-case scenario, if Broadwing pays the remaining $166.7m in principle and interest in stock, the current share count of about 68.8m (after payment of the second installment on 19 November) will dilute a further 46% to 100.7m shares—based on the recent share price of $5.50—after the final payment is made on 19 February 2006. If by that time investors award a successful Broadwing with a more reasonable enterprise-value-to-sales multiple of 2.0 versus today's 0.67, they will have tripled the stock price to nearly $18 despite the dilution. If the dilution ends earlier and/or revenues increase, then we can expect even more upside in the price.

CEPHEID (CPHD)

PARADIGM PLAY: MICROELECTRONIC MACHINES FOR DNA IDENTITY

NOVEMBER 19: 8.45, 52-WEEK RANGE: 6.16 - 13.56, MARKET CAP: 354.96M

A Tredennick favorite in the MEMS space, Cepheid is a leading candidate to create a lab on a chip for identifying dangerous toxins in an era of terrorism. Many analysts anticipate healthy growth in the biothreat market over the coming decade. For Cepheid, growth has been robust, with 3Q04 revenues up 25% sequentially to $14.1m and anticipated 2004 revenues of $49m up 163% over 2003.

The revenue increase is being driven largely by sales of the second-generation product, GeneXpert, which adds automated sampling to the first-generation SmartCycler that enables rapid genetic analysis of a sample. Though still in the final stages of development, GeneXpert is already being sold to the biothreat market through a Northrop Grumman led consortium developing an anthrax tester for the USPS. Sales to Northrop should continue into 2005. In addition, Cepheid shipped its first ASR (analyte specific reagent) primer and probe sets for identifying B. pertussis and HSV, and in the current quarter will begin shipping four additional ASR products which identify infectious disease organisms on the SmartCycler System.

Growing pains have been remarkably mild. Though net cash decreased sequentially from $54.5m to $38.3m, the erosion was due mostly to payment of upfront license fees and an additional $20.7m in fees to be paid over the next two years, both of which are accounted for in the figures. Otherwise, the company is burning a mere $2.5m to $3.0m per quarter and expects to turn cash flow positive soon. Although product mix is trending toward higher gross margin products, royalty and license costs should keep overall gross margins in the mid-40s through next year. Cepheid will also increase spending on R&D and manufacturing to accelerate delivery of clinical products, sacrificing near-term profits for long-term growth; the company project gross margins well into the 60s in five years.

Knee-jerk market reaction to the 2 November conference call (due possibly to the push-out of anticipated profits) caused the stock price to plummet 14% from $8.98 to $7.72 at market open on 3 November, presenting a brief buying opportunity which savvy investors apparently took advantage of—at market close the price had rebounded past the previous day's close to $9.14, yielding a sizeable enterprise value of 7x anticipated 2004 revenues. But Cepheid is a long-term player in a market that's in its infancy, and today's investors should plan on holding for at least five years. If revenue ramps at the same rate next year, we would be looking at an enterprise value of a modest 2.7x sales at today's market cap with still a long future ahead.

For investors, the questions are, How quickly will the biothreat and clinical testing markets grow, and how does Cepheid compare to the competition? Cepheid believes it has a unique product platform that performs more broadly across more markets than competitors such as Roche. Additionally, GeneXpert has a higher specificity and sensitivity than known rival products. Management expects continued growth in sales to the clinical market, and plans to expand its sales team accordingly. Clinical trials on the GeneXpert System are scheduled to begin this quarter and market launch the first half of next year. Expansion of sales into the life sciences market in China (currently less than 2% of company-wide sales) should be enhanced through a new distributor showing early signs of effectiveness.

EQUINIX (EQIX)

PARADIGM PLAY: WHERE STORAGE AND BANDWIDTH CONVERGE

NOVEMBER 19: 38.68, 52-WEEK RANGE: 22.00 - 41.38, MARKET CAP: 713.99M

Equinix booked a record-breaking 87 new customers in its third quarter and raked in additional orders from 47% of its existing customers, swelling its cabinet count to 10,750, up from 10,100 last quarter. Up also over last quarter is average recurring monthly revenue per cabinet, which is now $1,270. Third quarter revenues were $42.4m, a 37% increase over the same quarter last year, and recurring revenues were $40.1m, up 41% from last year. Customer cross-connects within EQIX business exchange centers also grew by 39%, over the same quarter last year.

With content giants Amazon.com, Walmart, and Netflix all planning to begin shipping DVDs over the Net rather than though the mail, we expect EQIX's momentum to continue. Where better to store the flix?

JDS Uniphase (JDSU)

PARADIGM PLAY: COMPONENTS GALORE FOR THE FIBERSPHERE

NOVEMBER 19: 3.14, 52-WEEK RANGE: 2.84 - 5.885, MARKET CAP: 4.53B

On the surface, JDSU had a stellar quarter, with revenues leaping ahead 11.5% sequentially to $194.5m, building on a crescendo of sales increases of 3.5% ... 5.8% ... 8.1% respectively since last September. The current surge was spurred by a 23.7% increase in sales of communications products, now 54.6% of total revenues. Long haul continued to come to life, growing to 25% of communications sales. In addition, sales of circuit packs grew by double digits in the quarter, and reconfigurable optical add/drop multiplexers (ROADMs) began shipping.

However, diving below the sunny surface, we find troubled waters. Despite double-digit revenue growth, gross margins slid from 25% to 22% over the past three quarters. Diving further down, gross margins get even worse. Correcting for the net benefit of $8m in sales of zero cost inventory (from previous inventory write-offs) this quarter and $3m last quarter, gross margins actually declined sequentially from 22.2% to 18.8%.

Trying to bring us back to the surface, management sees gross margin improvements once the transfer of manufacturing to China is completed. (Currently, JDSU employs over 2,000 workers in China.) Gross margins of 30% in the "mid-term" and 40% in the "long-term" were suggested. But in JDSU's world, product ramps never cease, and the company may always be playing catch up. For instance, five technology challenges are still blocking the necessary increase in yields in digital surface mirrors. While JDSU claims superior contrast compared to other mirrors in the industry, they are still in the "prototype" stage by their own admission, even as demand surges. Of the five production challenges, one or two seemed elusive to management.

Putting a positive spin on its forecast of a disappointing flat-to-down-8% December quarter, management treated the current quarter as a bump in the midst of a "mid-single digit" growth environment. In view of the sales crescendo cited in our first paragraph, we would have felt much better had management instead called the next quarter the anomaly rather than the current one. Now, with the admission of limited visibility, anticipation of "continued" single-digit growth in future quarters seems more like guesswork, and the forecast of gross margins holding in the low 20s only increases the roil. More hopeful is the claim that FTTH could increase revenue "significantly" in the second half of calendar 2005 and in 2006.

An expensive stock, JDSU continues to trade at an enterprise value of close to 5x estimated calendar-year 2004 sales and 4.2x estimated sales for 2005 based on the 5% per quarter growth assumed by management. Though the company is awash in liquidity and continues to improve operations, at the current market price investors should await signs of more rapid growth and significant improvements in gross margin before buying more shares..

ZORAN (ZRAN)

PARADIGM PLAY: AIR KING—DSPs FOR DIGITAL CAMERAS & DVDs

NOVEMBER 19: 11.42, 52-WEEK RANGE: 9.48 - 22.48, MARKET CAP: 492.88M

As we anticipated last month based, based on Zoran's earnings warning, 3Q revenue grew sequentially at a more-modest-than-expected 15.5% to $119.7m; the company had originally anticipated sales of $130m for the quarter. In the two days following the October 5 preannouncement, share price fell over 13% to $14.10, just above the 52-week low of $13.40 reached on July 27. However, despite the correction, we warned that patience would be a virtue with Zoran. The company blamed its slowing growth on Chinese DVD-player manufacturers depleting excess inventory due to tightening credit lines in China. Since the inventory correction began late in the quarter, it looked as though it might continue into 2005. Not great news for Zoran, considering that it is one of the largest suppliers of chips to the DVD market; DVD revenues are 62% of company-wide sales and China accounts for 40% of Zoran's revenues.

Zoran confirmed our suspicions by issuing a massive correction beyond anything we would have predicted, forecasting a 37% revenue slide to $75.5m in 4Q—39% below the $124.2m anticipated as late as last summer. This ends Zoran's three-quarterly-run of revenue increases, though the 4Q04 sales forecast is still $9m over 4Q03 revenues and the anticipated calendar-year 2004 sales of $379.5m are still up a marked 75% over 2003 revenues of $216.4m.

Responding to the news, the stock price plummeted to a new 52-week low of $9.48, less than half of the 52-week high of $22.48 reached in January. The stock is trading at an enterprise value of 1x the 2004 revenue forecast, down only slightly from the 1.2x it was trading at last summer on significantly higher anticipated earnings. Thus, patience might be a virtue a just a bit longer on the assumption that the Chinese correction is not over. In addition, the company cited the beginning of a general inventory correction across all its markets.

With digital signal processors, graphics ASICs, and full-motion camera chips spread across the consumer electronics landscape—Zoran's long-term prospects should remain positive. The balance sheet is strong and the company's DVD market-share continues to grow during the downturn. Hence, future dips may present buying opportunities.

points are carbon atoms and that you have a sheet of these atoms and their connections. Roll the sheet into a tube so that the atoms on one edge of the tube make connections with the atoms on the opposite edge. You've made a carbon nanotube. This tube has two measures that determine its useful properties: the tube diameter and the angle of sheet roll-up. Depending on the angle of roll-up, the nanotube will be either a metallic nanowire or a semiconductor.

According to their proponents, carbon nanotubes are an industrial panacea. They can serve as a circuit's interconnecting wires, as its semiconductor logic elements, or as its storage cells. With a tensile strength more than a hundred times that of stainless steel, nanotubes can be cables or textile fibers. The end of the carbon nanotube can be a strong electron emitter, which means that "lawns" of vertical nanotubes would make excellent emitters for displays. A carbon nanotube also makes a great tip for the probe on an atomic force microscope (AFM), because of its durability and its small diameter.

Following a familiar path, the suppliers of instruments, microscopes, materials handlers, and experimental equipment will be the first to make money from the nanotech industry. That's happening now. Second to profit are suppliers of bulk materials, such as the nanoparticles for sunscreen lotions, for stain-free and wrinkle-free clothing, for fingernail polish, and for coatings on lenses and other surfaces. As the production of nanomaterials comes down the learning curve, these businesses will grow. Next come simple components such as AFM tips and storage elements.

All-optical networks will always be superior to optoelectronic networks in both latency and reliability

It's much harder to build single-walled carbon nanotubes, with the desired structure, and to place them or to grow them where they are wanted in an integrated circuit, than it is to work with bulk materials. While bulk-materials applications proliferate, controlled applications in integrated devices will grow slowly.

For investors the key problem is that most of the companies with new ideas and new applications in nanotechnology (IBM, NEC, HP) are either so large that nanotechnology applications will have a miniscule influence on their growth, or they are private startups that are not suitable for investment (from a risk or from an opportunity perspective).

Such realities are obscured by the usual market projections depicting a current commercial market for nanoscale materials, tools, and devices of \$7.6B in 2003 and estimated growth to \$28.7B by 2008 (a 30.6% CAGR). Merrill Lynch publishes an index of 25 nanotech companies. Punk, Ziegel & Company also has a nanotech index. Four of the top companies on the Merrill Lynch list are materials companies: **Amcol International** (ACO), **Applied Films** (AFCO), **Altair Nanotechnologies** (ALTI), and **Symyx Technologies** (SMMX). Three of the top companies are instrumentation or semiconductor capital equipment companies: **Caliper Life Sciences** (CALP), **Ultratech** (UTEK), and **Veeco Instruments** (VECO). Among public companies, market capitalization of the materials and instrument providers dominates. There are two biotech companies, **Biosante Pharmaceuticals** (BPA) and **SkyePharma** (SKYE), and no semiconductor companies. We like **Cepheid** (CPHD) (see page 4), a biotech chip company that uses DNA to identify biothreats, and Microvision (MVIS) (see July 2004 *GTR*) a MEMS (microelectromechanical systems) heads up display innovator with a nanotech subsidiary.

The mix of companies will change slowly with time; instrumentation has been around a while and bulk materials are the simplest nanotechnology applications. While there are many applications for nanomaterials the most interesting applications will be in traditional semiconductors. Today's transistors are so small that they do not turn off completely. As transistors get even smaller, leakage currents increase rapidly. Over time, carbon nanotubes offer the promise of replacing today's transistors with precision-built, leakproof transistors.

Follow the red brick path

The International Technology Roadmap for Semiconductors, commonly referred to as the ITRS, is developed by industry participants, not outsiders. It is the participants' best guess of how semiconductors will progress in the next fifteen years. It is updated annually; the 2004 update edition should be posted on the ITRS web site (*http://public.itrs.net*) by mid-December. The roadmap is a detailed look at the issues and answers for the next seven years and a general look at the issues for the following eight years. The roadmap identifies the problems that must be solved by the industry to reach succeeding fabrication process "nodes" or targets.

Semiconductor process problems for which there are no known solutions are sometimes called "red brick walls." The semiconductor industry has a history of overcoming these problems. The industry is always approaching several red brick walls, but it has yet to be stopped by one. One wall crumbles and the next becomes visible. Here are some red brick walls facing the industry today.

First, at three atomic layers, gate oxide layers are too thin to continue to scale or to make high-quality transistors. One molecule more or less at *any position* in the layer varies the thickness by 25 percent to 33 percent!

Second, leakage currents are too high. In today's high-end microprocessor designs, leakage currents take up *half* the power budget. And, in two process generations, leakage currents rise tenfold.

Third, processes no longer scale linearly. Once upon a time, dimensions, voltage, and power shrunk with predictable relationships. Those relationships are breaking down. Oxide thickness is at its practical limit. Voltages, too, are near the minimum needed for the circuit's transistors to work. If the gate oxide thickness and the voltage cannot scale with other dimensions, then the speed and the power of the circuit likewise will not scale as desired.

Fourth, the semiconductor industry's top-down methods of etching patterns into bulk materials are nearing an end. Feature sizes are half the dimension of the wavelength that is being used to define the patterns. That cannot continue much longer. Alternatives are exceedingly expensive and are costly to develop.

Fifth, the cost of semiconductor processing equipment approximately doubles with each generation. The amortized cost of this equipment, over the number of chips produced, now rivals current prices of the chips.

Sixth, characteristics of on-chip transistors are no longer uniform across a chip. To make polysilicon conduct electricity, semiconductor manufacturers "dope" the wafer with atomic impurities. When transistors had large gate area, wafer-level doping provided uniform distribution of the impurities among gates. But, as the gate area shrinks, some gates are heavily doped and some gates are lightly doped. The result is statistical variation of transistor characteristics. Circuit designs, which for decades have counted on collective variation of on-chip transistor characteristics, can no longer do so. This statistical variation among the chip's transistors will be important for processes below 90 nm.

I'm no semiconductor-processing specialist, but these are the problems I can think of without doing research. There may be more problems, but this list is scary enough. The lists of five and ten years ago probably looked equally serious, yet the problems were overcome. Carbon nanotubes could be a major part of the solution.

Micro metamorphosis

The PC has grown from essentially nothing at the introduction of the IBM Personal Computer in 1981 to an expected volume of 160 million units in 2004. At its peak, it consumed 40 percent of the dollar value of semiconductor components. It was the first "killer app" for semiconductors. Engineering emphasis is now moving from tethered systems to mobile systems, which have provided a second killer app in the form of the cell phone with its hundreds of millions of annual unit sales. Emerging mobile applications have shifted the design emphasis from cost performance to cost-performance-per-watt. Divining the next killer app for continued semiconductor-industry growth, pundits now look to nanotech.

I suggest a different view. Instead of thinking of the PC and the cell phone as killer apps, think of them as successful proofs-of-concept. Under way is a microchip revolution that rivals the importance and the effects of the industrial revolution. We don't need a "killer app." What we are seeing is the transformation of entire industry segments, as semiconductors carry intelligence into the aging products of the industrial revolution. From automobile drive trains to medical prosthetics, mechanical and electromechanical systems are giving way to the intelligent electronic aids of the information age.

The automotive industry is in transition from analog to digital, from mechanical to electrical, and from isolated to connected (to the rest of the world via the Internet). The film and video industry is in transition from analog to digital and from isolated to connected. The consumer-products industry is in transition from analog to digital, from tethered to mobile, and from isolated to connected. (The consumer-products industry also offers high-growth opportunities in supplying standard appliances to emerging economies.) The biomedical industry is in transition from analog to digital and from wet laboratories to bioinformatics. The telecom industry is in transition from copper to mobile, from copper to fiber, and from electrical to optical. The computer industry is in transition from desktop to embedded.

These transitions will transform these industries. Think of photo processing, for example. It is moving decisively from analog film and wet-chemical development to digital storage and color printers. Transitions in a host of industries will create large markets for MEMS, for electrical and electronic components, and for computers and software for many years ahead. The first big transitions, from machine-room computers to personal computers and from wired telephones to cell phones, once viewed as "killer apps," were really the pioneering applications that opened the floodgates for the general transition from analog to digital and from dumb to intelligent systems.

An information-age automobile will use its own sight, sound, and motion sensors to assess its situation and it will use active suspension control, braking, engine control, steering, warnings, and other active, intelligent measures to prevent accidents and loss of control.

Today, people sit *between* the real world and the computer. People feed information to the computer and people interpret the computer's results. In the next phase, people will move to a supervisory role. The computer will collect and analyze its own information. The computer will manage its own sensors and actuators. The computer will decide what information to collect, how to collect it, how to analyze it, and what to do with it.

In a 90-nm semiconductor process, 100,000 transistors fit on a small grain of sand. But for the device to be useful, it must communicate with the outside world. Setting the size and power requirements of the system will be sensors and actuators and the transistors that drive and detect their interactions and bridge the scale to ordinary transistors. There is a fixed size and power cost in interface sensors, actuators, and transistors and a variable cost in computational transistors. Make the sensor too small and its samples will not be representative; make it too large and it will use too much energy and it may waste valuable materials necessary to the sampling process. Make an actuator too small and it cannot create noticeable effects for human interaction; make it too large and it will use too much energy.

In MEMS, the market is growing, but units are growing at the same rate, which means the price per unit is not falling (as the prices of semiconductors do). The price is in the packaging. New devices enter at higher price points. Revenues were $5B for 2003 and are projected to be $5.4B in 2005. The top 30 MEMS manufacturers have 60 percent of the market.

Key markets are optical switches, RF switches and relays, and actuators. The best-established markets are accelerometers, inkjet heads, and blood-pressure monitors.

Survival of the smallest

Biological systems, through billions of years of evolution, have developed ingenious, efficient means of locomotion, sensing, information storage, self-assembly, self-repair, and redundancy. Natural systems operate at low energy levels and often use efficient, ambient-temperature chemical reactions. We are now able not only to investigate and marvel at these implementations, but to borrow ideas from them. That's where nanotechnology and bioinformatics come in; they are the extension of electronics to the molecular domain. We will get atomic-scale devices through nanotechnology and we will borrow ideas from biological systems.

Nanomagnetics is a startup that borrows a solution from nature. The ferritin molecule is the protein for iron storage in humans. Nanomagnetics replaces the iron in the molecule's interior with a cobalt-platinum alloy and then spreads the material on a surface to create very dense magnetic storage. Similarly, today's memory components, SRAM, DRAM, and flash memory, are inadequate for mobile, intelligent systems. Market demand will foster development of new memory with the speed of SRAM, the density of DRAM, and the non-volatility of flash memory.

Axon Technology is a candidate, but there are many others as well, such as **Coatue, HP** (HPQ), **Molecular Electronics Corp.**, **Nanosys, Nantero, Rolltronix**, and **ZettaCore**. Nanosys is working with Intel and Nantero is working with **LSI Logic** (LSI). The eventual winner may be based on molecular transport, on carbon nanotubes, or on something entirely new. But it will arrive; in fact, several may arrive.

All these gains—nanotechnology, bioinformatics, and MEMS—are enabled by progress in semiconductor electronics; they are derived from progress in semiconductor electronics; and they are natural extensions to semiconductor electronics.

Even though today's transistors are good enough for most applications, the semiconductor industry, far from being mature or moribund, is at the beginning of a long period of growth. We are at the leading edge of a massive digital transformation as intelligent systems displace mechanical and electromechanical systems. Humans will move to a supervisory role as intelligent systems independently collect and analyze data. As a part of this intelligent revolution, nanotechnology's biggest successes will be in semiconductor applications, probably first in memory chips.

The PC and the cell phone have given us the confidence to trade places with the computer. We will allow the computer to become the collector and analyzer of data. Our willingness to put computers between us and the real world will enable widespread transformation of industries.

– Nick Tredennick and Brion Shimamoto
November 18, 2004

Got Questions?

Visit our subscriber-only discussion forum, the Telecosm Lounge, with George Gilder and Nick Tredennick, on www.gildertech.com

GILDER TECHNOLOGY REPORT PUBLISHED BY GILDER PUBLISHING, LLC AND FORBES INC. • Copyright ©2004, by Gilder Publishing, LLC

291A MAIN STREET, GREAT BARRINGTON, MA 01230, TEL: (413) 644-2100, FAX: (413) 644-2123 • EMAIL: INFO@GILDER.COM

EDITOR IN CHIEF	EDITORS	EXECUTIVE EDITOR	TECHNOLOGY ANALYST	MANAGING EDITOR	ART DIRECTOR	SUBSCRIPTION DIRECTOR	CUSTOMER SERVICE
George Gilder	Nick Tredennick Brion Shimamoto	Bret Swanson	Charlie Burger	Mary Collins	Peter Johnstone	Rosaline Fernandes	Sandy Fleischmann

For subscription information telephone toll free: (800) 292-4380 • WWW.GILDERTECH.COM