THE 7TH ANNUAL

TELECOSM CONFERENCE

August 25-26, 2003
The Resort at Squaw Creek
Lake Tahoe

Hosted by George Gilder
& Steve Forbes

Gilder
Forbes | telecosm
CONFERENCE 2003

THE 7TH ANNUAL

TELECOSM CONFERENCE

August 25-26, 2003
The Resort at Squaw Creek
Lake Tahoe

Hosted by George Gilder
& Steve Forbes

C O N F E R E N C E T R A N S C R I P T

Contents

Welcome Remarks

monday

august 25

8:30 am

George Gilder

Chairman, Gilder Publishing, LLC
Editor in Chief, *Gilder Technology Report*

Welcome to Telecosm. It's great to have such a full room and to welcome a lot of new people. It's the Turnaround Telecosm, but I also sometimes think it might be "I told you so" Telecosm.

There's been too much talk about bubbles in recent years, and I believe we should start by recognizing that the Telecosm companies managed a 9,000-fold rise in Internet traffic over seven years, and that was a heroic achievement—using the old equipment that was available in 1995 would have cost around $39 trillion. So this was a heroic achievement in many ways. The 9,000-fold rise in traffic was ignited by an approximate hundred-fold rise in 1995 and 1996. In some ways this was a miracle, not a bubble. I believe that this achievement was indeed accomplished, and it was largely destroyed or crippled by policy errors of savage deflation, which destroyed debtors around the world at the very time that all the telecom companies incurred huge debt in order to meet this transformation and explosion of traffic. I believe that ultimately what occurred follows Wriston's law: capital goes where it is welcome and stays where it is well treated. I don't know whether Richard Vigilante invented this law or not; he wrote the book in which the law was propounded, or ghosted it with Wriston. Wriston's law is close to the Telecosm these days because not only has the book *Telecosm* been published in all the Asian languages, in China, Korea, and everywhere else, and is doing far better in Asia than in the United States, but also $3 or $4 trillion of market cap and capital and technology was removed from the American Telecosm, where it all originated, and shifted to Asia. Today, Korea, a country with half the per capita income of the U.S., has twenty times the bandwidth per capita and five times the bandwidth absolutely to homes and businesses. It has 11 million households connected at an average bandwidth of about 5 megabits per second, a million households attached with VDSL between 13 and 20 megabits per second, and a massive program underway today to get another 2 million

Capital goes where it is welcome and stays where it is well treated.

households connected with VDSL at 50 megabits per second.

Korea has 11 million households connected—about 73 percent of the total households with real multi-megabit pipes. In the U.S. we have about 20 million households connected—about 20 percent of our total households with a meager form of broadband. The U.S. had a terrible depression in the Telecosm and a dot-com crash, and Korea is running most of its economy through the Net. One-third of Korea's GDP is comprised of online transactions; 70 percent of its stock trades occur on the Net; 50 percent of its banking transactions are on the Net; there's an ever-expanding proportion of retailing and teleconferencing galore. Its Hanoro telecom start-up of about three or four years ago has laid far more bandwidth to homes than any American telco, including the heroic sponsors of Telecosm (Verizon), who, I hope, will take a cue.

There's also a company in Italy called eBiz that has 70 percent of the households in Milan fibered and 50 percent of the households in Rome, and for $70 a month they'll give you 10 megabits per second. They are in over six Italian cities and expanding to eight or nine Italian cities. It's really mortifying that all these technologies, almost all of them originated, developed, financed, expanded in the United States, have been deployed fully outside of the U.S. There must be some reason for that.

What I've always said is the key problem is that we re-regulated telecom in 1996. We didn't intend to re-regulate it, but a million words of deregulation were readily translated by the lawyers into many more millions of words of litigation that in essence paralyzed our local loop. From that point of view, it's depressing to read that the FCC recently issued another 275,000 words of new regulation in their triennial review. It's just an incredible maze of new regulations, really, applying in 51 different jurisdictions, 51 states, each one with a different public utility commission, and two markets per state. Even Rhode Island isn't permitted to be a single market. There must be two markets separately regulated for the dollar threat of monopoly. I think this is a fundamental misconception about capitalism that's pervasive in Washington.

A million words of deregulation were readily translated by the lawyers into many more millions of words of litigation that in essence paralyzed our local loop.

Washington believes that capitalism is a desperately dangerous system that's likely to careen off into monopoly at any moment and that these monopolies are deeply disruptive and that it's worthwhile to cause general wreckage and rubble across the economy in order to prevent someone from possibly having fun or making money. This is a serious problem, but at the same time the news from Asia, from Italy, and from other countries around the

world is thrilling. It's a thrilling opportunity for the Telecosm because what we learned again in Korea is that when you deploy broadband pervasively what you get is another approximately hundred-fold rise in Internet traffic in a little over three years. And this hundred-fold rise that Korea has just undergone recapitulates the hundred-fold rise that ignited the great Telecosm boom back in 1995 and 1996. I hope this is a portent that as broadband is rolled out, and it will be rolled out in the U.S., we will have another nonlinear upsurge of traffic that will take most of the country and most of the experts by surprise and will precipitate a kind of panic deployment and panic buying that corresponds to the depression that we have recently undergone. This is a moment for the turnaround. It signaled in Asia but it's a tsunami that's going to sweep around the world and will even come back to the United States, and I just hope that this time we do it right.

To discuss some of the economic dimensions of this transformation that's underway, one of the critical policy errors that was committed was the deflation marked by the 25-40 percent increase in the value of the dollar against all gold and commodities and other currencies. In the beginning in 1996, this deflation punished debtors around the world, all debtors denominated in dollars, and it turned out that a great many of those debtors were companies deploying vast new Telecosm technologies and infrastructures. The problem began with economic errors, and I think it's fitting to open this event with a discussion of the economy, and we couldn't have a better person to do so than Andy Kessler.

Andy is full of surprises. He spent five years at Bell Labs designing chips, and most chip designers don't explode into the financial community, but he did. At Morgan Stanley he followed in the footsteps of Ben Rosen, whose footsteps I also trod as a newsletter writer. Andy later started his own fund, Velocity Capital, which was a tremendous success in the late nineties. And he was smart enough to get out of that and become a writer. His first book was *Wall Street Meat*, and many of you have read it, I suppose. It's an absolutely fantastic book—hilarious, brilliant, insightful, a terrific read. He also has a big serious book about the future of the American Empire on the way, and he's becoming an intellectual—another surprise. But today he's going to perform on the economy.

Washington believes that capitalism is a desperately dangerous system that's likely to careen off into monopoly at any moment.

You heard me. Give me 200 shares of SBUX.

Practicing what he preaches, Intel's Paul Otellini logs on via Wi-Fi (compliments of CloudX).

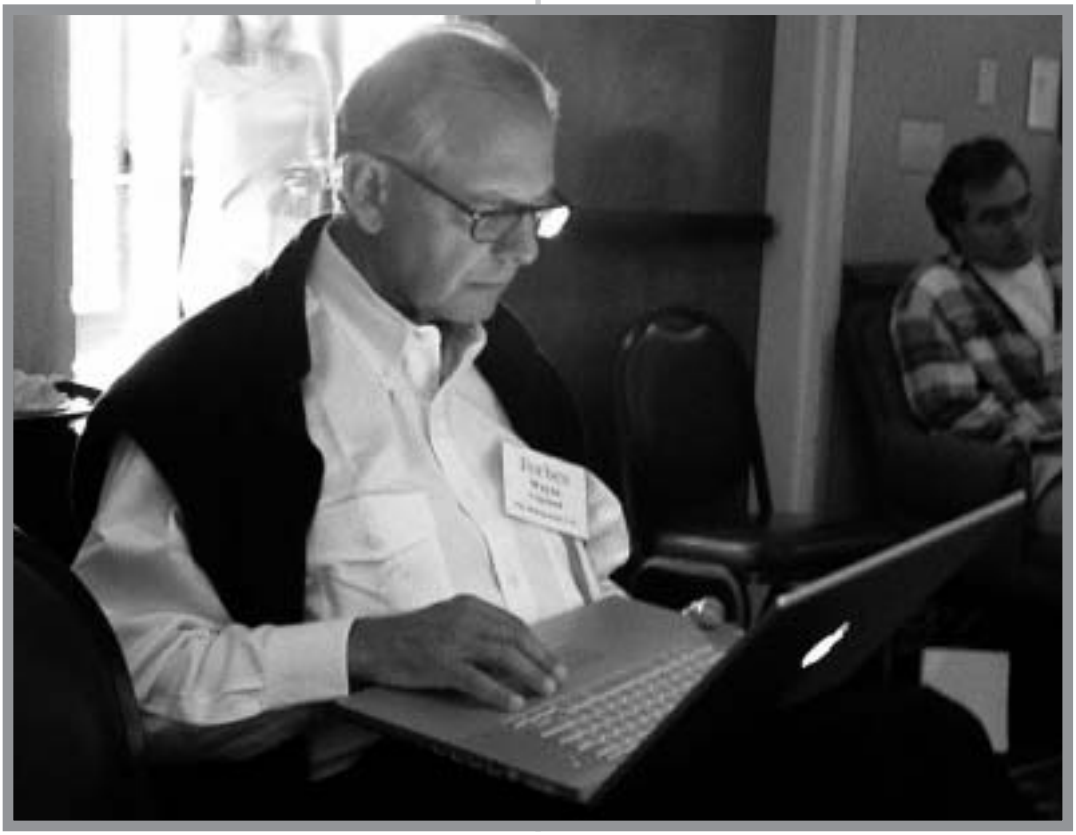

The ghost of Telecosms past hovers behind Wayne Copeland looking for a heads-up on the next 10 bagger.

Buffett, Bogle, and Dividend Plays

An economy can't run on dividends and index funds alone. Says who? The Great Debate.

monday
august 25
8:40 am

Andy Kessler:*

It's that intellectual thing that's becoming a problem—either you are or you aren't. I worked on Wall Street for probably a few too many years, and of course if you tell people that today, they snicker behind your back and very often right in front of your face. A lot of people think Wall Street is a clubby place or that the stock market is a deck that's stacked against them. I never thought that until recently when I bought the stacked deck. Sure enough, you go into the clubs, there's the Jack of Clubs, Jack Grubman; the Ten of Clubs, Frank Quattrone; the Four of Clubs, Mary Meeker; and the Seven of Clubs, Henry Blodget. Yes, Wall Street is a clubby place. These are all characters with whom I've had the fun of working and are part of the book (*Wall Street Meat*) that George mentioned. Did someone write down those little quotes? I need more blurbs for the book. But one card is missing. It's the Joker. I think Warren Buffett is the Joker, but we'll get to that soon enough.

Blaming Wall Street analysts for your losing money in the Bubble is like blaming your Members Only jacket or leisure suit for not getting dates in the '80s. You can't do that. You only have yourself to blame. Investors only have themselves to blame. There were structural problems, we know that now looking back, but the stock market is so important. It is the mechanism that allocates capital in our economy. It's all about access to capital. The stock market magically funds innovation and stifles capital from those dead companies that are dragging us down or holding us back. But we all know it doesn't always work that way. The Bubble, if we still want to call it that, overfunded innovation, and now the pendulum has swung the other way; perhaps it's starting to head back and funding innovation

* Former hedge fund manager • Author, *Wall Street Meat* and *Running Money*

again, we'll see, but today we're stuck with an era where there are new heroes. Warren Buffett is a new hero. John Bogle and his index funds are new heroes. Dividends are the new heroes. And to me that unholy trio is the anti-Christ of innovation.

What I want to do is wind a short tale that perhaps lends some insight into how to stamp out those three evils. Because if you want to look forward, it helps to look backward. I went back in time and tried to find some Silicon Valley models of innovation that lend some insight into the future. And I found one. He's a guy named John Wilkinson. I don't know if anyone invested in his company, probably not. John Wilkinson is the Ironmaster of Shropshire. It sounds like it's out of *Lord of the Rings*, but John Wilkinson had iron foundries. In 1775 he had this huge problem. King George laid on John Wilkinson a huge order for cannons. There was some issue in Lexington and Concord, or something in April of 1775, and the British needed some cannons to go *boom* and take out those nasty Minutemen. John Wilkinson was just the guy to give the order to because John Wilkinson was an ironmaster. At that time, to be an ironmaster you took iron ore and mixed it with charcoal; you laid it together, ran bellows to heat up the charcoal, and then melted the iron out of the iron ore; you poured it out into a cast, which was the smelting process. The problem with using charcoal to smelt iron ore was that it was organic. You burned off as much of the organic stuff as you could, but you were left with a lot of sulfur. When you mixed it with the iron, the iron that came out with iron ore was very brittle. So they were looking for new ways to process the iron. To make those cannons, John Wilkinson had a proprietary, barrel-boring tool—a lathe. He was not only the ironmaster; he was the machine-tool master. So he'd pour that iron into a solid cast and used the lathe to bore out the cannon. The windage—the gap between the cannon ball and the edge of the barrel—was so fine that with a little gunpowder, *boom*, you'd take out those Minutemen.

Blaming Wall Street analysts for your losing money in the Bubble is like blaming your leisure suit for not getting dates in the '80s.

Unfortunately, charcoal was no longer the way to go after a while, and a guy named Abraham Darby—there was actually a whole family of Abraham Darbys—who came up with another process. I don't know if they used a

mass spectrometer or not, but they figured out that coal, or sweet pit coal, had a lot less impurities and that if you used it to smelt iron, you'd get much more rigid iron out of it.

In the past, foundries were by rivers because you used a water wheel to pump your bellows and to heat the process, but Wilkinson moved up to the hills. In the hills were the mines for the coal and iron, and since he didn't need wood, he was fine. Believe it or not, this is all going to wind up to the present and show you why Warren Buffett is such a buffoon. So Wilkinson's up in the hills. And what did he lose? He lost his power source. He didn't have the water wheel anymore. And he had this huge order for cannons, and the materials were backing up, and people were yelling at him, and so what he did do? He had horses. He used a bunch of them to go round and round in a circle. I don't know if they went clockwise or counterclockwise. They'd stare at each other's tail and go around and around, but it wasn't enough to pump the bellows to get the sweet pit coke hot enough. So Wilkinson looked around for another power source. It was hard to miss actually. It was a very crude steam engine used to pump out the flooded mines. Which brings us to a company called Boulton and Watt. You know James Watt; you may not know Boulton. Boulton was Watt's venture capitalist, and Boulton owned two-thirds of Boulton and Watt; Watt only owned one-third. That's how it goes, right?

Believe it or not, this is all going to wind up to the present and show you why Warren Buffett is such a buffoon.

Anyway, Watt was a great inventor, and he had a Professor Black at Glasgow University who was working on latent heat, and he figured out that you can create a vacuum by having an external condenser that you cooled down. Anyway, Watt's steam engine was sitting there pumping the water out of Wilkinson's mines and it was full RPM. It could lift two tons of water 165 feet. Wilkinson said, "Great! This is going to run my bellows." He went over, hooked it up to his bellows, and it didn't work. Big problem. The steam engine was a piece of garbage because it was basically a 1- to 3-horsepower engine. Wilkinson, like any good engineer, took it apart. There was this six-foot-high cylinder that the piston was in where a vacuum would be created, and unlike his cannon barrels, the windage was all over the place. It was very difficult to create a vacuum. He used wet hemp, which I think he got from Jamaica, to seal the barrel. So he thought, "Wait a second. I've got this proprietary boring tool." And Wilkinson created his own cylinder with a very tight windage, and sure enough, he took the Boulton & Watt steam engine and just by fixing that barrel turned it from a 3-hp engine to about a 25-hp engine, and *bang*, *flash*, the British had a hundred-year empire.

It's funny, but that's how these things work. There are two guys doing completely separate things, and all of a sudden they come together; there's a flash of innovation; there is more wealth created. How did it work after that? Well, not only did he make his cannons—which is modestly interesting because he lost that work since he didn't fill the order till after Yorktown—but there were a bunch of guys doing interesting things with textiles. In the past manufacturers would take the cotton that was grown down South and spin it and twist it. But cotton is a terribly itchy material. It was like wool. It was cheaper than wool because you grew it

rather than fed it to generate it, but it was really scratchy. Up until about 1775, unless you were royalty and wore silk, you had itchy underwear. But in effect, the British Empire was about making comfortable underwear. You can thank Wilkinson and Watt for that.

There was another guy named Nathaniel Crompton who had this tool called the spinning mule, and instead of just spinning and twisting the cotton it would take it and stretch it five feet. When you stretched it five feet, you made it smooth. It took a lot of power to do that. There were mills with spinning mules by rivers for the water wheels; unfortunately only when there was a rainfall was there enough power to stretch the cotton. If you hooked up a Boulton & Watt with a cool, new Wilkinson cylinder, that was it. The British had this triangle trade, and everyone in the world traded their grain and corn and gold for comfortable underwear. The Silicon Valley part of the story is not only was it comfortable underwear, but it was cheaper than you could do it at home. There was an elasticity of that comfortable underwear by which some spinster, sorry to call anyone in France or Germany or anywhere else on the continent a spinster, couldn't make it as cheaply spinning at home as the British could by stamping it out with their steam engines and mills. There are a zillion lessons in that, but did anyone figure out where Warren Buffett was in that story? Warren would have funded the farmers or the horse-feed companies. He wouldn't have gotten anywhere close to the innovation, so, of course, he's the horse's ass.

One last thing that's important for later is that Boulton & Watt had a twenty-five-year patent on the steam engine, and Wilkinson had the proprietary rights to sell the cylinders for them. Patent law was very important and so were intellectual property rights because they didn't allow others to copy. Patent law was quite strong. Nor did they allow their machine to be exported out of Britain. But rather than sit there and milk the profits by selling the steam engines, they put steam engines into factories and charged one-third of the annual costs of the horses they replaced. If you had twenty-five horses at 10 pounds each, you would pay 80 pounds a year to Boulton & Watt for the twenty-five-year duration. Rather than pay dividends to Boulton who traveled the world, for instance, Watt kept inventing. He invented the implanted gearing, the double-action steam engine, the tilting forge hammer, all sorts of things. He kept investing back into the company.

So what's my beef with Warren Buffett? Well, Warren Buffett is about the status quo. He's an investing genius; I'll give him that. I'm not taking anything away from him, but if you dig a little closer, he's an investing genius for himself, not for you, not for the benefit of our country, not for the benefit of society. Basically, Warren Buffett runs an insurance scam. Insurance scams are quite simple to pull off: you collect premiums for something and then you avoid paying out any claims on them to some future date and hope to invest well enough ahead of paying out claims so that you generate profits. Martin Frankel—I don't know if anyone knows him—he had the same insurance scam except he invested in Greenwich real estate and women with loose morals.

Buffet's an investing genius for himself, not for you, not for the benefit of our country, not for the benefit of society.

We could spend a whole hour on why insurance is a scam, but I don't really mean to do that. What I want to get across to you today is that Warren Buffett is not my friend, your friend, anyone's friend. He's Warren Buffett's friend. He has a cult following of people who believe in the Oracle of Omaha. But let's go through all the ways he hates your guts. He's going to hate my guts after I'm done with this talk.

First of all, he's got this high stock price, and he refuses to split his stock. As I last checked, his stock was $76,250. If you're making thirty grand a year, you're not really investing alongside Warren Buffett, which is what everyone thinks you're allowed to do. And if you do—he actually did come up with B shares of one-thirtieth, which is now $2,542—you don't have any voting rights with those shares. You're somewhat of a second-class citizen if you can afford even the $2,542, and when you buy those shares Wall Street will clip you with a nice couple-hundred-dollar commission.

Warren Buffett is also for limiting CEO pay, and he doesn't like paying stock options. He thinks options are an expense. What that means is that it's an honor to work for Warren Buffett. I don't know how he attracts any of the best and the brightest if he doesn't provide them with the incentive of stock options. If he wants to do that and milk the profits from his company, that's fine. But I don't think he has to push that way of thinking on the rest of us. He also promotes unproductive companies. This might be controversial, but one of his investments is Fruit of the Loom, which we now know from my story is a circa 1775 tale. Making even cheaper underwear is not what moves society along, nor does Dairy Queen, although I do like Dairy Queen.

Buffett is also anti-technology. His quote is "I just don't get it. I've never invested in technology companies because I just don't get it." Instead, he was a big investor in the *Washington Post*; he's still a big investor in the *Buffalo News*, but like the old line about the railroads that forgot they were in the transportation business, newspapers are in the business of connecting people. Today, if you add up the market cap of Dow Jones, The New York Times, and the Washington Post companies, it comes to about $15 billion. And eBay, which is in the same business of connecting people and nothing more than a glorified classified advertising business, has a $35 billion market cap. His pronouncements telling people, "Don't invest in that innovative tech stuff; it's just going to get you in trouble" has actually gotten him in trouble. He could have dumped all his newspaper shares and invested in the same business.

Warren Buffett is for corporate governance, and who besides Andy Fastow isn't for corporate governance? I looked up his board of directors. See if you find a similarity here. Warren Buffett is on the board, great. Susan Buffett is on the board, great. Howard Buffett is on the board. Charlie Munger is a partner, of

course. Ron Olson, Charlie Munger's law firm partner is on the board. Walter Scott, who was at Peter Kiewit and actually did some co-investments on the side with Berkshire Hathaway, is on the board; Thomas Murphy of ABC recently joined the board—of course Warren Buffett was a big investor in ABC. In other words, he might say he'd like there to be corporate governance everywhere else but for himself.

He likes to keep his investments secret. In fact it just came out recently that he had a couple of investments for which he didn't want to file what are called 13F filings. He can invest, but unlike mutual funds he doesn't want to have to tell the rest of the world. Of course they slapped him down on that. He's against the elimination of inheritance taxes. He's rich but everyone else who's born has to hit the reboot button and start again. And finally, he's for population control. Now in 1798 Thomas Maltheus proposed that population grows geometrically and food grows arithmetically, and I think that's been proven wrong. So Warren Buffett is not someone who should be hero-worshipped, not someone to be followed for his teachings.

Another guy who bubbles my iron ore is John Bogle. John Bogle started a company called Vanguard, which invented index funds. There was a problem in the stock market where there were all these mutual fund managers, and less than 25 percent of them would outperform the market in general. They would charge high fees for managing your money and then screw up. John Bogle said that he could fix that. Rather than try to beat the market, why don't you just "be the market?" It's like Danny Noonan in *Caddyshack*: "Be the ball, Danny." Bogle just says, "Be the market." He would charge 18 basis points rather than 75 basis points that JP Morgan might charge institutions, and voilà, it worked.

I'm all for cutting management fees, now that I'm no longer in the money management business, but the entire market is now indexed. Everybody is indexed to the S&P, and that promotes mediocrity. It's all about mediocrity: its companies are sticky. You may not want to own IBM or Schering-Plough, but when you're indexing the market, even if you're in Fidelity Magellan, you're basically in an index fund. It's not labeled that way, but these guys are all closet indexes. And this country and its stock markets should not be about mediocrity. I'm not adamant about very much, but I am very adamantly against mediocrity.

It's like Danny Noonan in Caddyshack: "Be the ball, Danny." Bogle just says, "Be the market."

Finally, speaking of mediocrity,

there has been a recent dust-up over the idea of the double taxation of dividends. Now I'm all for eliminating taxes on anything involved with the stock market—capital gains, dividends—that's fine. But I don't think one type of payout of the stock market should be favored over another. The problem with dividends is that when companies pay out a dividend they're telling you as an investor that they have nowhere better to put the money. They can't invest in themselves. And if they can't invest in themselves, I don't want to invest in them. All I'm saying about dividends is don't use them as an investment criterion. It's nice that they pay it out; it's nice that taxes have been halved; let's make the tax rate zero on it. But dividends are a bribe by under-performing companies to keep you in their stock. If they didn't pay a dividend, no one would own their stock and they would lose their access to capital. My message is unless every company paid out every penny of their earnings as a dividend don't look at the dividend line as an investment criterion. Kodak pays a 6.3 percent dividend; the stock has been going down all year. Schering-Plough went from 60 to 15, and then Friday they announce: "Ugh. Business is terrible. We have to cut our dividend." "Cut my what?" Anyway, that is my message. Innovation is key.

The next two days are about finding and listening to lots of innovative companies that are going to move the market and society along and generate wealth for all of us. There are some Wilkinsons and Watts and, I hope, maybe even some Boultons out there who are increasingly easy to find, so don't let the hero worship of Buffett, Bogle, and dividends sway your investing decisions.

George Gilder:

I first heard of John Rutledge way back in 1981, I believe, when I was working with Congressman David Stockman. Stockman had discovered stocks, and he kept talking about stocks and flows. He really babbled about it, and I tried to follow his argument and found it very elusive. But he was very smart and very quick, and he told me that to get to the bottom of stocks and flows I had to talk to this guy John Rutledge. Stockman, at the time, was forming Reagan's economic policy and Stockman's key advisor during this period, the guy he really respected most, was John Rutledge, an economist, of all things, but who later went into a number of other fields, including the management of money, with great success. He, like Andy Kessler, is full of surprises. He's now exploring physics as a source of analogies for the economy. He'll give a different perspective on markets and their uses.

John Rutledge:*

Thank you, George. I remember in 1981 that both of us had a lot more energy. It was a hoot having a boss who was a cowboy. I remember President Reagan's first staff meeting. He said, "I hate taxes, I hate the Russians, and I hate inflation. Work something up." And that's where the Reagan plan came from, folks. The first time I saw the Reagan plan it was one sheet of yellow 8 1/2 x 11 paper in Alan Greenspan's pocket, who'd been traveling around on the campaign trail. When we finished the plan, it was the "Rosy Scenario," and we really had a lot of fun

* *Forbes* columnist • Partner, Rutledge Capital

with Rosy at that time. In those days the one thing to remember was that inflation was 15, interest rates were 20, the marginal tax rate was 70 percent, and the USSR was a country. All those things have changed.

I totally agree with Andy on innovation. I have to say that a year ago and again today talking about physics in front of Carver Mead is really an embarrassing thing to do, but if you'll promise not to stand up and scream, I'll move ahead. There's a wonderful man who died about two months ago. His name was Ilya Prigogine and Ilya Prigogine wrote about a thing called "far from equilibrium physics" and irreversible change. Pretty fascinating things. I'm going to talk about that just for a minute. A little book of his you ought to read is called *The End of Certainty*. It's medium-simple, not too geeky. Prigogine's basic pitch is that the universe is a far from equilibrium place, and everything we know is a far from equilibrium thing. It's not the equilibriums from the textbooks. He said that the irreversible nuclear reactions inside the sun maintain our ecosystem by maintaining essentially price differentials or heat differentials across the universe. That's what makes energy and entropy flow. It's what makes, as Forrest Gump said, "shit happen." And that's what life is all about: it's about these energy flows. It's what allows life to happen. It's also what allows entrepreneurial activity to happen, as Andy was talking about earlier. To me there are energy flows in economics. Anything that makes more of them is better than anything that makes less of them. If there are more of them, and they're actually well aimed at something that's going to create value, that's even better. So I'm a fan of innovation and I'm a fan of capital gains cuts and I'm a fan of almost any tax cut.

I remember President Reagan's first staff meeting. He said, "I hate taxes, I hate the Russians, and I hate inflation. Work something up."

I'm a guy with attention deficit disorder, so I have about forty jobs. One of them is owning and building companies. Another one is working for governments. This last year I've worked for the White House on both the dividend

tax cut, where I'm for it instead of against it, and I'm one of the guys on the Iraq rebuilding task force, just because I spent a lot of time in the Gulf and I know how to do business with a desert camo on my head.

A friend of mine said that this is the least intellectual administration they've ever seen. President Bush, when you look in his eyes, is a guy who does things because he believes them not because he thinks them. To him the dividend tax cut was very simple. It was an equity thing; it had to do with double taxation. It might be right; it might be wrong; it doesn't matter. It was actually his bugle. He decided on it and from then on it was, "Do you want it?" or "Do you not want to have it?" And I said I'll take it and every single tax cut when I can get it, because taxes are an impediment to energy flows and taxes make for less economic activity. I've traveled 15-million miles in the last thirty years, and I've seen rich people and poor people, and I prefer being rich. The case for low taxes is the case for liking to be rich.

The Fed rate cuts have all been cosmetic, and they haven't really accomplished much of anything. The soft spot is on Alan's head, and it's time for Alan to go.

I'm going to give you a couple minutes of tax-deductible slides. *Supra videri* is the favorite expression of a wonderful guy, Leonardo Da Vinci. Well, wonderful might be stretching it. He actually was tried twice for pederasty in Florence, but that's another story. He thought the secret was knowing how to look at something. I'm going to get you past that. I'm going to tell you about a couple of things: something about stuff, something about seven years, and something about the tax cut.

Actually I came here from the Aspen Summit, which is for telecom regulators and telecom guys, and it's a sucking-up event. Somebody asked me what's the difference between a growth company and a mature company? Growth companies criticize the government; mature companies suck up to the government. That's the one sure way to know. Anyway, regulators were thick on the ground in Aspen. We had the chairmen of pretty much everything in one hotel. They were worrying about the best way to regulate competition, which is an oxymoron if I've ever heard of one, so I suggested not looking at prices because price controls can lower prices and destroy things. Don't look at numbers of competitors. Look at capital stock. Anything that makes the capital stock bigger will make us richer. If you can run regulatory policy in a manner that maximizes the capital stock, you will be a good regulator. How do you do that? For most of the things regulators do, we'd tell them to get out of the way. But there are decent things that happen, like setting the foot and the yard. Standards setting, things like that, but that's about where it ends.

There's a great myth going on out there, which is that 9/11 caused a recession and Alan Greenspan saved our butts. Then we had a soft spot and now the economy is back on track and there's a great recovery out there. Almost all of that is wrong: 9/11 did not cause a recession. It was already really terrible before that. The Fed rate cuts have all been cosmetic, and they haven't really accomplished much of anything. The soft spot is on Alan's head, and it's time for Alan to go.

We actually have only half a recovery, and it isn't really very good out there. The economy is medium-lousy. I'll talk about why in a minute. It has to do with

deflation, which George was talking about. Don't get excited. Most people write about deflation like it's herpes or Argentine loans. Either you have it or you don't. It isn't like that. Deflation can be more or less, and it can be different for each company and each person. There are a lot of prices out there: they can all go up or down. When a lot of them go down, it's very painful.

I want to tell you a story of two recessions. One of them is this bubble everybody talks about. I object to the idea about talking too much investment or too much capital. That's impossible. You can have the wrong capital for the work you've got to do, but you can't have too much of the stuff. Capital is what makes you rich, but there's a capital spending issue and there's a working capital issue.

There are two recessions that have been running side by side in this recession trench. The first one is almost over; the second one isn't. We know about this Bubble thing everyone writes about. There was new law in the 1990s. It invented new industries; CLECs are an example of that. It also did things like drill holes in your street and run cable and the like. Twenty percent revenue growth was the norm when everybody raised money. Three-quarters of the money raised in the late nineties was for these few industries. The 20 percent pro-forma revenue growth turned out to be 15, 10, 5, 0, and negative 10 percent, and as the top-line growth faded so did the quality of those loans. In the fourth quarter of 2000, the Fed began to worry about that. They called up their pals at the controller in the Treasury Department and said you should do something about these banks because the controller regulates the banks. The controller dispatched the goon squad out to the hinterlands, and they visited every bank and said that they needed to lower their exposure to commercial industrial loans. Guess what? If you loan the money to somebody who eats it for dinner, it's a venture deal. If you pay someone to drill a hole in your street and shove a fiber down it, you're not going to get any of the money back because it's gone. The only thing you can hope for is enterprise value of whatever they built, but if it's shut down, it's gone. So the banks could not get any money back from their customers they loaned to. Therefore the pressure from the controller was deflected to small, dumb, and ugly companies—companies that produce a product and have inventories, receivables, and need working capital to meet payroll. This whole issue was accidentally deflected onto a second story, which I'll tell you about in just a minute. In the meantime, the bankruptcy auctions fetched returns of five, six, seven, eight, nine cents on a dollar's worth of book value. In those auctions, the product went on junk dealers' shelves. The junk dealers sold it for twenty cents on the dollar and doubled their money. However, twenty cents on the dollar is less than a hundred cents on the dollar, so if you needed to buy something in a box over the last three years, you bought it from a junk dealer, not from Cisco or Nortel.

The largest seller of IT equipment last year was IBM. Junk dealers were the number two seller. They dominated the resale market, and the nice thing about it is that when the boxes are gone, they're gone. And they are gone! That happened, it's over, and now we're back to business again. That's why you saw an IT-

If you needed to buy something in a box over the last three years, you bought it from a junk dealer.

pickup in spending last April over the last twelve months because the junk dealers are no longer in business.

Don't get too excited about a recovery. Every time people do, they buy things they shouldn't, and they've been wrong six times in a row.

The other thing is that the product in the boxes becomes obsolete. If the company is clever, at the very moment there's a bankruptcy auction and the merchandise gets into the hands of the bankers, the company changes the code and makes it obsolete. Then the next guy who buys it from the junk dealer has to buy it from you instead. There's an obsolescence issue here as well. That, too, is over. Consumers are now buying software; they're not buying systems; they're not buying machines; they're not buying buildings; they're not buying plants and equipment. Investment spending is abnormally biased toward software right now, which is unsustainable. Don't get too excited about a recovery. Every time people do, they buy things they shouldn't, and they've been wrong six times in a row. They're going to be wrong this time too. There's half a recovery, not a whole one.

The second recession I mentioned, bad loans, regulators, hyperventilated goon squads, bank closings, working capital. C&I loans, the money you buy for payroll, was $1,104 billion in the fourth quarter of 2000. This week [August 25, 2003] it's $925 billion. So $180 billion worth of inventories and receivables has been sucked out of businesses in the last two and a half years. That number fell last week too. It's fallen for about 150 weeks in a row, and the economy can't recover until it turns back up again. I have numbers on that if you'd like to see them. The myth is that the Fed has fixed everything. The reality is all the Fed did was lower interest rates for public companies because private companies can't get money from public markets anyway. They've got to get money from a bank or from their uncle. This is true monetary policy. We can talk about that later. This is really an arteriosclerosis story, and we have to find a way to break through it.

The technology bubble ended some time ago. There's some capital spending

happening, but these are the numbers on C&I loans. Right up until today they continue to fall. This is the cause of the deflation that George was talking about. Not enough working capital. One of the ironies of it is that when companies don't have working capital, the customer can't get product so they buy it from another source as an import, and the product becomes pregnant with working capital. In other words, the Central Bank of Japan or China or Germany provides that company with working capital, so it's a way of evading the credit restraints we have, making the trade deficit part of this as well.

Monetary base has finally started growing. It's actually what the Fed owns: it's their balance sheet, but there's a dirty little secret. It is that for most of last year currency holdings were rising very strongly. Here's the metaphor: I'm Alan Greenspan, and George, you're going to be the guy with the Treasury bill. I'm going to buy your Treasury bill and do an open-market operation. To make the monetary base grow, George has two choices. One is that he can stick the T-bill in his jeans. The other choice is he can deposit it in the bank. If he sticks it in his jeans, that's the end of the story. If he puts it in the bank, the bank has more reserves and can lend it to someone, and it goes around and around and around. For most of last year, currency holdings went up so fast—currency is 90 percent of the monetary base—that none of it ended up back in the banks again. Why did that happen? Monetary policy was hijacked either by Enron or Martha Stewart or Saddam Hussein or by someone who scared the hell out of people so they put their money in their jeans. But the good thing is since we invaded Iraq, currency holdings have gone flat. When currency holdings go flat, bank reserves rise and they continue to rise. If you want to be an optimist, people unclenching their buttocks means they hold less currency and put their money in the bank; the banks eventually begin to lend it out, and that's when the recession will end. But that's a premature story.

The reality of all this is that we have half a recovery now, but long term we have a bigger problem. It's a wonderful problem in a sense that it's technology growth, because it's what makes productivity grow. Each person is more productive than before. But guess what? If you own a business and every employee is 5 percent more productive next year, you've got 5 percent more product to sell. If no one will buy it, you've got to fire someone. The job of the Fed is to make demand grow faster than productivity so that companies will hire employees to work in their businesses. The Fed has not been doing that and we've got to find a way to get them back on track.

The biggest worry investors have today is not being invested. Staying invested in the stock market is the moral of the story.

There are two main ideas I want to leave you with. The first is that an investment model portfolio theory is totally wrong. Instead, there are ways of measuring the actual business risk of things relating to performance. The second is that any time you see stories in the newspaper about macroeconomics that say, "Who's spending the money anyway?"—throw them away. What you should be interested in is network failure, system failure. The blackout that happened on the East Coast [August 2003] is exactly the same as this credit crunch I'm describing. It's a cascading network failure where something goes

wrong in one place and it ripples down. The reason why it happens in the case of the Fed is that there's one node to that network that's connected to everyone, and that's the controller of the currency. This is not a very robust network because of central controls. If you replaced macroeconomics with Prigogine's far from equal equilibrium dynamics, irreversible thermodynamics—which is about collisions, complexity, and system failures—you're going to have a much more interesting view of the economy. Credit expansions and credit system collapses are macroeconomics actions and making people productive and accumulating capital was what you wanted to do as a trend.

I'll be around for a couple of days if anyone has any economics questions. I'd be glad to talk to you about them. The story is half a recovery. Which means 2 to 3 percent growth, not 4, 5, 6 percent recovery growth. Falling prices to flat.Which means falling to flat employment. Which means the Fed is benign. Which means bond rates bought too low have popped back up. There's still some risk in the bond market; the stock market remains 20 percent or so undervalued; the dividend tax cut is worth 10 percent on the stock market in its initial impact on dividend payers and 20 percent more on people who don't pay, but eventually will pay dividends; it's a lowering of cost to capital. If you stay invested in this market and pick good places that are going to grow, it's a great time to be invested. The biggest worry investors have today is not being invested. People who lived through the downfall and then sat out the upturn have paid the price, have paid the tuition, and are not going to get the rewards. Staying invested in the stock market is the moral of the story.

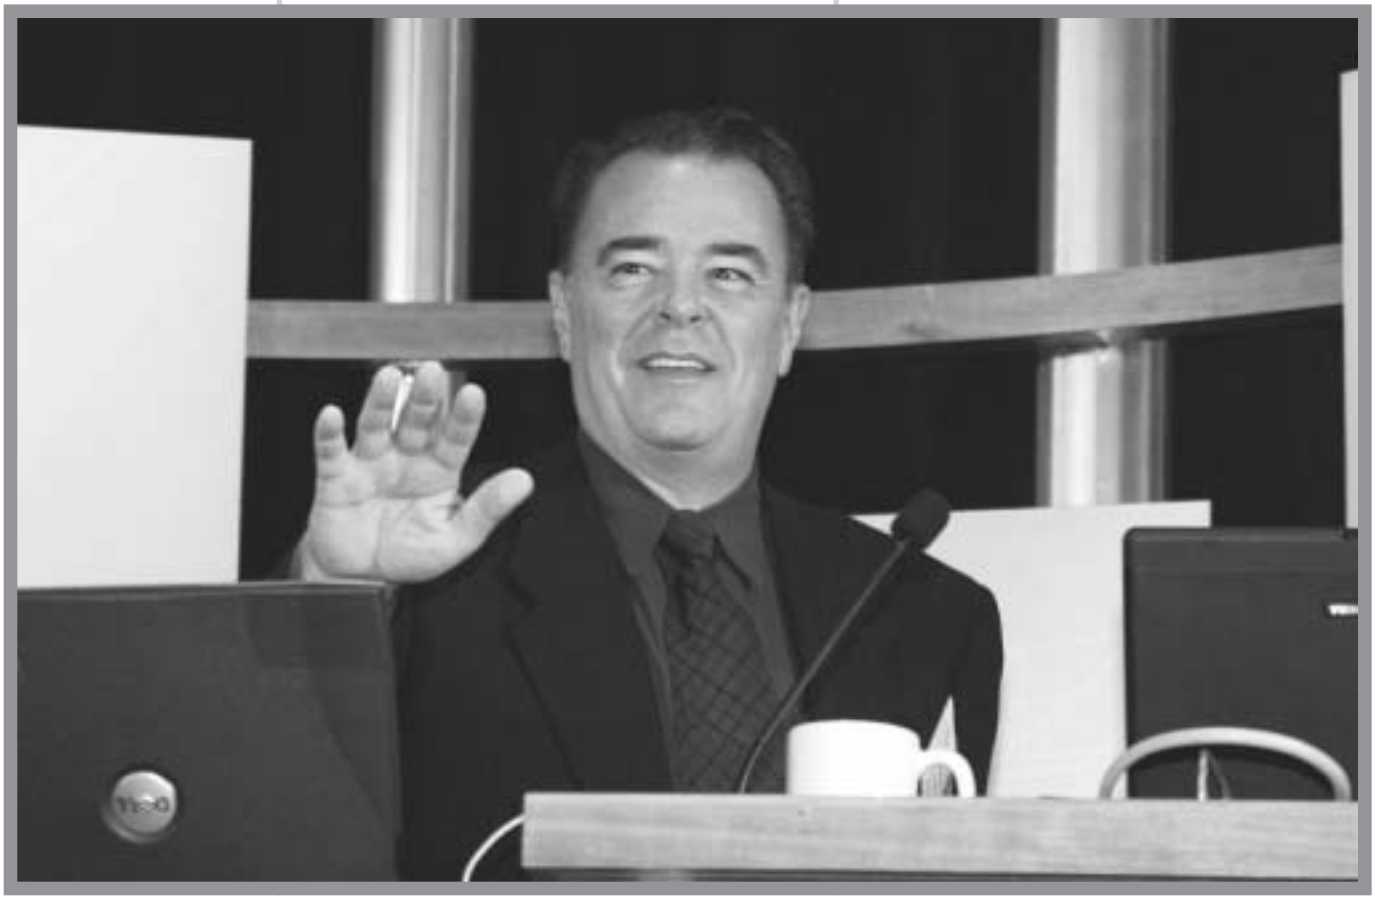

Boy, the service here is fast!

Saving a seat for George who forgot the conference starts at 8:00 am, not 8:30.

Is he just really tall or is that a miniature chair?

That's Ashby. "Always on" and "always fast."

BREW

changing the way the world looks at wireless

monday
august 25
9:30 am

George Gilder:

My favorite company for a long, long time has been Qualcomm. I discovered it some time longer ago than I can even recall, and I was immediately struck by the brilliant minds that were aggregated in the company. This is what really impresses me about a company. You meet people like Qualcomm co-founder Klein Gilhausen when you visit. It is an incandescent place, and it led me to a great assurance that they were going to triumph in their pursuit of CDMA. Now a further point about Qualcomm, which is a little-known secret of its success, is nepotism. Most of the time nepotism is destructive, but I think that if it's as well managed as Adam Bellow expounds in his new book, nepotism can be good, and certainly Dr. Paul Jacobs, with his twenty-five patents, his key contributions to CDMA, his management of the handset division, and general spearhead of Qualcomm for decades is a prime example of Qualcomm's intellectual qualities and of the redemptive properties of nepotism.

Paul Jacobs:*

It doesn't happen often, but that kind of an intro does leave me speechless. Being from an innovative company, at least I think we're innovative, that's paying dividends, I have to at least address the fact that I see dividends a little bit differently. I look at it as we're there working for the shareholders, and in some ways paying a dividend is somewhat like paying a paycheck back to the owners of the business. If that's a bribe, I guess Qualcomm is paying me a bribe every day to come to work when they're paying my salary. I think there's no problem in reconciling innovation and dividend paying. We're generating an awful lot of cash with our businesses, and we're simultaneously

*** President, Qualcomm; Executive Vice President, Qualcomm Wireless and Internet Group**

investing in R&D, and I think we're doing a lot of really interesting new things. I don't actually buy the argument that they conflict.

There is one point that I want to make before I get started. Today in San Diego, we're celebrating the fifteenth birthday of OmniTRACS, which is our messaging system for long-haul trucking. It's the system that we started first and it actually paid for the development of CDMA. The really cool thing is that for a long time it was seen as a very mature, low-growth kind of business. But CDMA is now feeding back into OmniTRACS. What we're able to do is take all the expertise we've built having this end-to-end data services business that was primarily a satellite-based system, and now we're running these things over terrestrial data networks and opening it up to all sorts of things besides trucks. We have everything from tracking construction equipment, which is very similar, to something that seems much further away, monitoring people's cardiac events on wireless heart monitors.

It's interesting to follow Qualcomm's history. We started out with a data services business, and then we built a fundamental data-carrying (and obviously voice is important) capability. Now we're able to take that data service business—and you'll really see more of this in the future—and drive a lot with that service concept, which goes to my talk here.

The wireless industry is in the middle of two transitions. One of them gets talked about a lot. It's the transition to 3G—a transition to new technology. But there's a related transition that's less visible, the transition from voice to data-enabled services. That to me is a much more fundamental change in the mindset of the carriers. It's not exactly related to 3G because a lot of these services are coming out on 2G networks, but obviously 3G will accelerate those capabilities. Voice began as the killer app of wireless, but it's non-voice applications that are driving things now. I'm not saying the revenues that are generated by the data services are going to overtake the voice services right away. In fact, the most aggressive guys are saying maybe 25 percent of the revenues will come from data in 2005. What's interesting is that marginal decisions of customers going into stores and buying phones are now being driven by non-voice products. Voice is commoditized; it's expected. You now buy a phone with a color screen and more memory. Coming to the United States strongly is picture messaging, a phone with a camera. It's those non-voice kinds of things that are driving what people are doing in wireless. That's a very, very fundamental shift. If you deal with the carriers at all, you recognize that they've had to retool themselves. They have to sell differently now. There are all sorts of applications available. Some are in the vanguard in Korea or Japan or in Asia, where things started earlier, but are coming to the U.S.

The transition from voice to data-enabled services to me is a much more fundamental change in the mindset of the carriers.

The U.S. is where there has been a very fundamental change, particularly on Verizon's network, where the ability was launched to download applications to handsets. There are tons of applications available, and every one of those applications was done by a third-party developer. We've brought a whole new community of people with their creativity and their innovation into the wireless industry, and that's really going to

help drive the industry forward. That's another part of the transition that we've seen of moving beyond voice. And you know what? It's really helping the carriers' economics. If we go to Japan and we look at KDDI, their net profit tripled in six months after they launched 3G. The great thing about what they did is that they did not sell it as 3G, even in supposedly data-crazy Japan. Because camera phones were extremely popular over there, they took a camera phone and instead of having the same resolution as their competitor, they used the fact that they could send more data over the air to have a higher resolution camera. They went out to the consumers and said: "Buy our phone. It takes better pictures." That's how you sell 3G, and that's how they've gotten these kinds of results. They're getting ready to launch the next generation of service, which is our EvDO technology that does peek data rates of 2.4 megabits per second and average data rates in the hundreds. We're seeing 800 kilobits per second now in San Diego. Their data speed lead is going to continue to persist. They've led in terms of the number of subscribers added, and they have a bunch of new and interesting services. One of these services is launching GPS-assisted.

The Japanese said, "Buy our phone. It takes better pictures." That to me is how you sell 3G.

The fundamental thing is that these guys are making money. SK Telecom, in Korea is doing a similar thing. They have launched our EvDO technology across Korea already. A million people are using it to download multimedia. The impact on their economics is that the average revenue per data user is up to over fourteen bucks. They're downloading broadcast TV. They're doing video on demand, and they're doing channels. People are very interested in buying this capability. The cost that they're selling this data over the air for is really important. The cost per bit had to be driven down because carriers are used to selling you voice for ten cents a minute at 4 kilobits per second data rates, or SMS messages for a quarter that are a few hundred bytes. To start doing multimedia over the air and do it effectively at a price people can afford, we had to drive the cost per bit down. And that's what we're doing. When you talk about where we are innovating as a company, we've been driving the cost per bit down for a long time. It's a key focus of ours and something you should watch relative to other

technologies that are brought to market. In any case, money is being made on a new service.

KT Freetel is in a similar situation. Comparing revenues between their 2G and 3G phones, there are about a thousand WANs in the 2G case, and over 10,000 WANs with a similar kind of multimedia service over the same EvDO technology. The interesting thing is to see the rate going from 1,000 to 10,000, and it's not because they're getting all the early adopters. These guys have penetrated their user bases pretty well. The $14 that I talked about at SKT was on 10 percent of their installed customer base. It's not just an early adopter phenomenon. You can see that when they went from their 2G to their first 3G, they got a 6x to 7x bump, and then they went to an even higher data rate and were able to provide newer services on top, which generated even more revenue. That's the key.

Look at the competitors, the pricing of the GPRS networks—the supposed competing technology, although it's half or less the data rate—is much higher. The fact that we've been able to drive the cost per bit down enables carriers to provide flat rate data services on the wide area. According to Verizon, for $79 you can get 1x data, which is up to 44 kilobits and average data rates in the 50-70 range. I downloaded 40 megabytes yesterday at my house over the 1x network at the same flat rate. We're in the position in the CDMA community to continue to drive down the cost per bit. We're already starting to see these services in Asia, and we're beginning to see the vanguard of them coming into the U.S. What that's led to is a great sense of adoption. Yes, GSM has a larger installed base. There are more carriers running it, but if you look at the growth rates, CDMA is winning. The spread of CDMA is going to be much more evident when all the GSM operators move to W-CDMA, because they need to get the capacity gains and the data rates to enable them to provide voice services and all the interesting data services out there.

Killer app is an overused term, and I apologize for it, but I don't think you or anyone else ever thought they'd see software retailing as a killer app.

Killer app is an overused term, and I apologize for it, but I don't think you or anyone else ever thought they'd see software retailing as a killer app. It's about providing choice to the end users to make the decision about what they want to optimize their handsets for. I don't know if any of you have used "Get It Now" on Verizon or on any of the other systems out there, but essentially we made the killer into a retailer. We built a retail distribution chain for software where the software vendors supply into a channel and re-aggregate: we take a wholesale margin on it; we supply that to the operators; the operators decide what software they want to have on their virtual shelf; they set the retail pricing; and the consumers buy it right on their handset. It's a one-click, downloadable kind of thing. We made it very simple for people to use, and there isn't a lot of advertising on it.

Fifty percent of people who buy one of those enabled phones on Verizon's network actively use this technology to download applications. What you've all been told about Americans not wanting to access the Internet on their phones because they're used to the Internet on their desktops is wrong. Consumers in the U.S. actively download applications. The key is that we didn't just take something from the wired Internet and shove it onto a small-screen phone. These

were applications tailored to the device. What a surprise. If you tailor applications to the device, the experience becomes more compelling to the consumer. The results are that there are almost 11 million handsets in use now. Nine operators so far have launched a service. It's a simple thing to integrate into a handset, and we designed it so it could be done on a very low-tier handset or a high-end handset. We also do it across technologies, by the way.

What we've all been told about American consumers not wanting to access the Internet on their phones is wrong.

You've all seen business plans with hockey sticks in them that say that this is how my business is going to take off, but it's not often that you see actual charts with hockey sticks in them. We're seeing this replicated over and over on all the carrier launches in terms of the number of downloads that are happening. Right now it's mostly CDMA operators, and we've got a couple of new ones in terms of customers, APPW in Taiwan and Reliance in India, which is going to be the largest telecom company in India very quickly. We're also doing trials now on GSM operators. Why do we do that? Because we want to drive more data usage in the GSM market so that they move to CDMA more rapidly. That's happening now. The way we do it is with a complete end-to-end system for developers. We put together a system that allows three guys in a garage to build an application and fit it into a global electronic marketplace, which the carriers can choose, so that the three guys in the garage don't have to travel to India to sell their application. When the user buys it, we facilitate paying the developers. The developers get a check from Qualcomm basically every quarter. This system is working. Developers in places like Brazil are building applications and selling them to Verizon in the U.S. In the mainstream PC industry the opportunities for small developers to become big is diminished because there are so many players, and it's difficult to get through the channel. But we've made a distribution channel that's very easy for small developers to get through. We're seeing handfuls of guys leaving video game companies, starting up something on their own, building applications, and then selling them on Verizon's network, the largest network operator in the U.S. Or on VIVO's network in Brazil or China Unicom's in China. If you're looking for investments, there are some very interesting things going on in this particular industry where people are building applications and are actually making money. There are real revenue screens being generated on the wireless Internet for developers right now. There's a huge range of brands that have gotten involved; there are big companies you know the names of, but particularly what

those big companies are doing is going to some other developer to write their applications for them. Then there are small guys you haven't heard of, and they're all in the process of building all sorts of applications. One of the nice things we're able to do that's important for the wireless industry is to bring these other brands into the industry.

One of the biggest problems the industry has seen, of course, is churn—consumers buying a handset on one network, moving to another network, and continuing to move along. If you can create a more lasting bond and not provide a commoditized voice service only, you can reduce problems and fundamentally affect economics. When you look at the CDMA carriers, they have better technology and better business models. We are already seeing that reflected in their economics, and we're going to continue to see that reflected in their economics going forward.

To give you some sense of what's going on around the world, Verizon is the largest CDMA carrier and the largest carrier in the U.S. right now. They've announced they're making $7.50 incremental average revenue per user on a mid-to-high 30s as their base. This is a big number for an operator to have, and that's all based on content. They've had lots of downloads. They also have one of those hockey stick curves in terms of the number of downloads that they've had.

What are the applications that are of interest? Everybody thinks it starts in Asia. Asia is very particular; they only like certain applications. Across the world the kinds of applications that are popular are fairly similar. One exception, in the U.S. is karaoke. You definitely have people in Japan and Korea sitting on trains with their cell phones using the karaoke app and singing to themselves. I don't think that would cut it here in the United States, but aside from that, I'm pretty sure everybody is going to get into casino. We all play solitaire on our PCs already. There's also content downloading of things like ring tones. Ring tones are a great 3G application. When we started out we didn't have data services at all, and they used to give you a sequence of keys to press on your phone and they would beep out the song. Now that we have medium-speed data services, you can download MIDI, which is a text description of the song, and it's actually synthesized back. It sounds better but it's kind of musak-like. KDDI launched a service in Japan where you download the actual waveform so the real signal is on your phone, and that's what's ringing. I apologize if it gets even more obnoxious, but I'm telling you that's one of the onsets of 3G. We're going to have even more obnoxious ring tones.

The only thing I haven't seen be popular in the U.S. is karaoke. In Japan and Korea people are doing the karaoke app.

One of the more interesting things about KDDI is that it's launched a position location system around the country. It's network-assisted and works indoors. They found that it takes a combination of position, location, and other things to make it compelling to consumers, so they have movie services, for example, that know where you are and tell you where the movies are playing around you. It even shows you the trailers. You can also take pictures and it marks the picture with the location of where you took the picture so that you can sort pictures by location. They're doing a lot of safety and security things as well. We've had a whole bunch of cases where stolen jewelry or lost children were found using these services. It's the combination of

different technologies that's the key reason the wireless Internet is different than the wired Internet. You have capabilities such as position sensitivity or time sensitivity, mobility or ubiquity, the fact that you always have it with you. We can take applications from the wired Internet and augment them. Chat is a great one. When you have a chat service on the wired Internet, your Buddy List tells you who's online. If you're wireless, you're always online. You can say, no, I don't want to be disturbed, but you're always accessible.

We can take applications from the wired Internet and augment them. Chat is a great one.

What's the next step? Maybe you want to know how many of your buddies are within two miles of you so you can all meet at the nearest Starbucks. It's combinations of attributes of the wireless Internet and the wired Internet that are going to build compelling services.

People think about Latin America as a place where maybe they're not willing to spend as much money on applications, but they, too, have the hockey stick chart in terms of growth. The numbers started out relatively small, but are growing very rapidly. Even in a very cost-sensitive market, applications are taking off.

China Unicom I mention because they just launched nationwide. They took Yao Ming, who is a well-known basketball star in China, and built contests around him. They took an outside brand, brought it in to build a service, and built contests where the winner gets to meet Yao Ming. All kinds of interesting things have opened up in the transition to a wireless Internet model that works. None of us would have expected that this is how China Unicom would be advertising their CDMA service. I certainly did not predict it, but it's actually based on a game that was written in China by a Chinese developer. It's being sold around the world and called "Prince of Persia," an adaptation of an old PC game. They say that 3G stands for games, girls, and gambling.

In the end the real key is that we're integrating a lot of new technologies. All kinds of features you've seen in your PC are being integrated into chipsets for phones. It's being done in a cost-effective, power-efficient, size-efficient way. Most of the world is going to get its access to the Internet through a handheld device. That's what's going to take the wireless industry beyond voice into a new stage. That's why Qualcomm has focused more on services going forward. There's a lot of opportunity when new markets open up globally and a lot of opportunity for investment. We're certainly investing, despite our dividend.

George Gilder:

Java has been one of my favorite technologies, and for years I've been predicting it would have a great use on handsets, and indeed it is having widespread use on handsets. I'm interested in your sorting out the role of Java in BREW applications. What kind of convergence is possible? What do you predict for these two competing technologies that also have complementary parities?

I think we made a mistake when we named the technology BREW. It sounds like it's directly competing with Java and we were thinking about a whole different drink.

Paul Jacobs:

Basically, I think we made a mistake in one sense when we named the technology BREW because it sounds like it's directly competing with Java. We were thinking about a whole different drink. The technologies are complementary in the sense that we can enhance the BREW system with any kind of extension, and a Java virtual machine is one extension. The way that we built the technology was to get away from having standards bodies. Any developer who wants to write in a certain language, be it Java or Flash or some proprietary language, the server knows when the application is downloaded and knows to download an extension to the phone to enable the phone if the phone isn't inherently enabled. We've seen a bunch of applications done that way. The one thing with Java that caused problems in the developer community is that it doesn't have an end-to-end system. People get paid differently through different carriers. The deals are all different and that's held back the community. It's a software layer that takes more processing power to run so the applications typically run slower than native applications, which we can do in BREW. I actually think that in the end Java is going to be used mostly by enterprise developers because that's what they're used to writing applications in. But people who are writing specific apps for phones to sell commercially, I think, are going to write mostly in C and C++ and write natively, just because they get that much better performance out of it.

The other thing that's interesting about Java is that Java security is based on the notion that content comes from everywhere. An executable program is just like a Web page. It comes down from anywhere. We don't know who the developer was so we can't really trust it. We need to use this piece of software to protect the device from the downloaded software. In the cell phone industry we know a lot about the device. We know about the user of the device, particularly because you're already making micro-payments on a cell phone system. Every phone call you make is a micro-payment. So we have a different attack on security. We force the developers to get a signature for their application, and once that application has left the developer, we know that that application hasn't been tampered with. We have a chain of responsibility all the way back to the developer. We have a financial model that's end-to-end, so if they cause a problem we'll revoke the app out of the phone, and there's a financial penalty associated with it. The reason why I think that's important is because I'm very concerned about all the Java phones that are out there with an embedded piece

The line for face-to-face time with Paul Jacobs grew so long we started serving coffee.

of software that's supposed to protect the phone, the device, and the network from a wide range of unknown applications that are going to come down and try to attack it. It's not like the PC where you can just do an upgrade of your software. If somebody cracks the security, the phones are going to have to be recalled in order to be fixed. The reason why it's complementary with BREW is that we can recall the virtual machine and upgrade that over the air. It's a key fundamental issue that you have the ability to upgrade things over the air, that you have this security model associated with it.

The wireless industry uses Moore's law in a different way than the PC industry does. We're throwing a bunch of it at battery life.

The wireless industry uses Moore's law in a different way than the PC industry does. We're not throwing all of Moore's law advantages at processing power. We're throwing a bunch of it at battery life, because I don't think anybody in the room would say that they like the battery life of their cell phone right now. You'd like it to last months. As we put in color screens and more applications that are running all the time, it's going to chew up battery power. We'd rather save the processing power, not go as high, and preserve it to get battery savings. That's why having lower apps that require less processing power like native apps is a good thing in wireless.

Wi-Fi Hot Spots

picking up steam or blowin' smoke?

Sky Dayton:*

I talked to Paul Jacobs just a second ago and he asked me why I wasn't wearing a tie, so I asked if I could borrow his, but I don't really know how to tie one. I thought it was apropos of the "debate" between Wi-Fi and 3G. I don't actually think there's a debate, and I hope today that we're going to be able to put that to rest. I also want to talk about how we build the roaming system for Wi-Fi, some of the challenges that Wi-Fi faces, and its true potential to help people get connected to the Internet faster and easier.

First, I wanted to talk about the key drivers behind Wi-Fi's success. Many of you are very familiar with all of this. Wi-Fi is incredibly fast. The air link is 11-million bits per second and up. The actual speed, however, depends upon the back haul and the location, so if we have a T1 that's connecting back to the

* CEO, Boingo Wireless

Internet from the Wi-Fi network that's in this room, the speed of the network would be 1.5 megabits. The Wi-Fi link itself will never be the limiting factor in the speed of your connection. Wi-Fi is short range, 100-500 feet, so it's ideal for concentrated environments. We have lots of people who want to access the Internet at the same time in a concentrated physical space.

For the first time in the history of the wireless business users already have the device that they need to connect to the network before the network is even built.

Wi-Fi uses free spectrum. In 802.11b and g there's about 83.5 megahertz of spectrum available, and in 802.11a there's over 300 megahertz available. There's a lot of talk at the FCC about opening up even more free spectrum, which is great. This is a huge amount of bandwidth. In comparison, a typical carrier using licensed spectrum might have 20 megahertz or less. Wi-Fi is also a global standard. In many ways it represents the TCP/IP of wireless. Everybody working around the world on an innovation helps everybody else, increasing returns. You don't have the bifurcation of innovation that you have in GSM and CDMA. The result has been massive competition and a very rapid commodization of the components that go into Wi-Fi equipment.

In 2002 there were about 20 million Wi-Fi chipsets shipped, and that should grow to about 40 million in 2003, about a hundred percent increase. That will be attended by a decrease in total revenue for Wi-Fi chip companies going from about 370 million to about 340 million. The reason this is happening is that chip sets are dropping so fast in price. In 2002 the cost was about $16. The prediction is about $8 by the end of 2003; $4 next year; and in 2006 people are talking about $2 chip sets. At this price, they can be put into pretty much anything, and I think this will happen.

Wi-Fi is incredibly cheap, driven by falling prices. A Wi-Fi card, something you might put into your laptop, was $700 three years ago. Today, it's basically free. If you buy any decent laptop, an 802.11 radio comes built in as standard. At the same time, the device that you might use to set up a network, which you plug into a wired Ethernet connection to your back haul and you power up—will broadcast Wi-Fi in a range of 100-300 hundred feet typ-

ically and cost about $2,000 three years ago. Today you can get them for under $50. Again, everyone will have them because the price is dropping so quickly. All these reductions in price have resulted in tens of millions of user devices with Wi-Fi radios out there today looking for networks to connect to. It's the first time in the history of the wireless business that users already have the device that they need to connect to the network before the network is even built. Unlike in the traditional carrier model where you spend lots of money for spectrum, billions of dollars in cap-ex to build out a network, hundreds of millions in op-ex to operate the network, and then pay a hundred to two hundred dollars for every device put into a user's hand, in Wi-Fi that piece of the economics just doesn't exist.

The other thing we're seeing as a result of the proliferation of Wi-Fi devices is that lots of networks are being set up today primarily in homes and offices. There are millions of residential and enterprise Wi-Fi networks today, and increasingly we're seeing the emergence of so-called "hot spots" where Wi-Fi networks are set up in public spaces. Hot spots initially serve people outside the home or the office who need to be connected. We've reached a point where dial-up isn't enough bandwidth. You could be sitting in an airport lounge, someone sends you a PowerPoint that you need, and you could miss your plane if you're using dial-up. There's a critical need for the kinds of bandwidth that Wi-Fi and other technologies can provide. There are 27-million business travelers in the United States who carry laptops. That's just the low-hanging fruit—people who need it right now. Dell says that by the end of this year all laptops will have Wi-Fi as a standard component. We're also seeing PDAs—HP, Toshiba, Palm—that have Wi-Fi, and we're going to see more. We're going to see Wi-Fi-capable cell phones; Motorola has already announced one. I believe that Wi-Fi will become the standard component in cell phones in the future. We're seeing Wi-Fi go into cars to serve telematic systems. You pull into your garage; your car docks with your home Wi-Fi network, uploads and downloads information. Think about Wi-Fi-enabled Game Boys. You're playing a game and you're connected to the Net and that device becomes something completely different. As chips continue to drop in price and because there's so much innovation on Wi-Fi around the battery profile and form factors, we're going to see it go into all kinds of things that we haven't even imagined. It's difficult to predict what that's going to be. I don't think any of us in 1993 would have predicted what the Internet has become and have been able to extrapolate from what's available today to what will happen in the future.

Hot spots are appearing all over the world. We've seen the emergence of a new kind of company.

Hot spots are appearing all over the world. We've seen the emergence of a new kind of company. At Boingo we call them "hot spot operators," or HSOs. These are folks who are setting up public networks in airports, hotels, and cafés. They are companies like Wayport, STSN, start-ups like Surf and Sip, major telecom companies like T-Mobile, AT&T, SBC, and Sprint. Pretty much every major carrier is in Wi-Fi in some way or another.

Verizon has taken an amazingly innovative approach by equipping payphones in New York City with Wi-Fi. I was in

New York just recently, and almost everywhere I went in Midtown I saw a signal called Verizon Wi-Fi. Next time you're in New York look for a small device that looks like a little black hat on the top of the Verizon phone booths. That's an omni-directional antenna. There's a DSL line coming into the payphone; they use the New York Power Authority loop that's already in place—it's really cool. No one is going to stand in a payphone and use their laptop, but they might use it across the street in a café or in an office upstairs. I was on the fiftieth floor of a hotel in Midtown and I saw the signal.

Target locations are high-traffic areas. Airports are the most important and the lowest-hanging fruit, followed by hotels and convention centers and then cafés and other public spaces. This business is characterized by a very low barrier to entry. It's inexpensive to get into it, but there is intense competition. Unlike previous attempts to build fast data networks to serve customers like Ricochet Networks, for example, the user already has the device they need to connect to the network before the network is built. It doesn't matter that there's no capital to build networks today. These hot spots are built in such tiny chunks that they fly under the radar of capital markets. For a number of reasons, Wi-Fi in a concentrated area provides a speed at a price per bit at a cost to deliver to the end user and thus a cost to the end user to buy that no other technology, 3G included, can touch. That doesn't mean that 3G isn't necessary; it's absolutely vital, but it does challenge the business case.

Today, we have about 5,000 hot spots live in the United States across the industry. We think that there will be a tipping point around 10,000 in the right places.

TeleAnalytics predicts 300,000 hot spots worldwide by the end of 2005. I think that's probably a year too optimistic, but we will see numbers like that toward the end of this decade. Gartner says there will be about 75 million hot spot users by the end of 2008, So this isn't just idle talk. In New York and San Francisco you can download Boingo software from our Web site to sniff out available Wi-Fi signals. Many are corporate networks that you won't be able to connect to because they're encrypted. Some of them are free networks. Some

are commercial public networks that you can pay to connect to.

There are four key barriers to adoption that we need to address as an industry. The first is lack of ubiquity: it's not yet in all the places you want to use it. When I travel I try to plan my travel to cities that have Wi-Fi-enabled airports and I stay in Wi-Fi-enabled hotels, but I don't think people should have to make that choice. That is not a good user experience. Today, we have about 5,000 hot spots live in the United States across the industry. We think there will be a tipping point around 10,000 in the right high-traffic places. When we think about ubiquity and Wi-Fi, we don't talk about square miles of coverage. That's just not the right organizing principle. It's about where people spend a lot of time in a concentrated physical space. If you were to multiply 10,000 hot spots times its actual coverage, it's a tiny number. It's less than a downtown area or a city, but that's not what matters. What matters is where people actually use it and whether we have coverage there. I think we're going to get to that point toward the end of 2003, beginning of 2004.

The second problem facing the industry is fragmentation. When you're in New York or San Francisco and you're walking down the street, you might pass a café, a hotel, or a bookstore with three completely different hot spots operated by three completely different companies, meaning that you would have to go and sign up with three different providers in order to connect. That's not going to work. It's like the early days of cellular, but much worse. Because Wi-Fi has this low barrier to entry in short range, there will be lots and lots of different companies out there building out the infrastructure. Unlike cellular, no one hot spot operator is going to control more than 10 percent of the total footprint. In cellular you could figure that 90 percent of your traffic was going to be on your own network and 10 percent would be roaming on someone else's. It's the reverse for Wi-Fi, so roaming is absolutely essential—something we address at Boingo.

We made a lot of strides in the last year in Wi-Fi, but like the early days of the Internet, you have to be somewhat of a geek to use it.

The third problem is ease of use. It's still just not easy enough. We made strides in the last year, but like the early days of the Internet, you have to be somewhat of

John Csapo of Samsung, Sky Dayton of Boingo Wireless, and Jeff Belk and Paul Jacobs of Qualcomm debate Wi-Fi vs. 3G.

a geek to use it. We have to make it a lot easier. We're making progress, but we have more to do.

Finally, for carriers that are getting into Wi-Fi there's a significant challenge in the Wild West aspect of the business. A user can roll into a city and see all these different signals and the carrier doesn't have any control over it—something carriers aren't prepared for. They're used to building networks that they then give their customers access to. This is about other people building networks. They build maybe a little bit of it, but then they give their customers access to all of it. How do they control that? That's something else we address at Boingo.

To help understand how it all fits together, we created a segmentation model for the industry that has four layers. Focus on just one of these layers is going to be critical to success. The first layer is the venue owners, the companies that own the physical real estate like your hotel chains, airports, cafés, gas stations—public spaces. This layer is characterized by lots and lots of fragmentation. The second layer is the hot spot operators—the companies that build the networks in each of those locations—like Wayport, STSN, the others. Lots of companies comprise this layer. There are over 200 hot spot operators around the world today. There will be thousands of them in the near future—big and small companies going out and setting up hot spots and lots of room for lots of competition. The fragmentation at the hot spot layer creates a need for the third layer, which is roaming. This is what Boingo does. We work with all the hot spot operators to create a single network that a customer can access. In turn we work with the major brands that provide that service to their end user: major carriers; ISPs like EarthLink and AOL; new types of remote access players like FiberLink that all want to market to the customer, bill the customer, and add value-added services. I think that this segmentation is going to become increasingly clear and important as people focus on one layer or another.

Wi-Fi is not a replacement for cellular; it's good at different things. It's a LAN not a WAN.

The Boingo system today has over 2,600 hot spots under contract, something that we'll be announcing publicly tomorrow. This is in 17 countries. It's not just a U.S. phenomenon. We have 30 airports with full or partial coverage, over 1,300 hotels, and over 700 cafés. We've really tried to focus on the key areas: airports, hotels, wherever business travelers go, where they really need this. We have 40 hot spot operators. The network is comprised of a lot of different people building out infrastructure, and it's growing very rapidly—over 100 percent growth so far this year (August 2003). I think that we can grow another 100 percent by the end of 2003. Our target is 5,000 by the end of the year, the largest aggregated roaming system for Wi-Fi.

I want to talk quickly about a myth I've been hearing and that is that hot spots can't make money. The reality is it's very easy to make money in hot spots. If you own a café and you want to set up a hot spot to serve your customers and charge for access, it's pretty straightforward. It will cost you, conservatively, about $150 a month to operate that hot spot. The numbers can be quite a bit lower, but you might spend $100 for a DSL line, $50 for depreciation of equipment and operating expense, or about $5 per day. Boingo will provide $2

in revenue per connection, so if you have 300 people walk in to your café each day a 2 percent utilization rate would yield you about $360 a month, or a 60 percent profit margin. The reality is as people show up with lots of Wi-Fi-capable devices, utilization rates are ultimately going to be much higher than that.

The second myth I want to address is that Wi-Fi will replace 3G. That's misguided and patently untrue. Wi-Fi is not a replacement for cellular; it's good at different things. It's a LAN not a WAN. 3G provides ubiquitous, relatively medium-speed, high-cost service, but it's going to be available wherever you are—in your car, out in a field, in suburbia. Wi-Fi is what you're going to use in key high-traffic hot spots. It's fast and inexpensive. Wi-Fi does challenge a business case for 3G in the sense that the highest traffic locations actually favor Wi-Fi. It's where Wi-Fi shines and where spectrum is the scarcest. Carriers are looking at that and saying that if we can offload our traffic onto Wi-Fi in these places where we really have spectrum scarcity, it's a great way to go. In the future hot spots will cover the highest traffic areas: homes and offices, clearly; private networks; airports; hotels; conventions centers; cafés. Wi-Fi challenges the original business case for 3G, but it doesn't replace it. 3G is critical for ubiquitous coverage and both will thrive. It's not a debate about either/or. It's both. The first companies that can integrate both for the end user are going to be huge winners.

It's not a debate about either/or. It's both.

The 3G Response

monday

august 25

10:35 am

George Gilder:

I'd like to introduce Jeff Belk of Qualcomm who's made a big hit on the Internet with the description of his attempt to travel through Wi-Fi, hot spot to hot spot, through Europe and other areas.

Jeff Belk:*

When George first positioned this panel, it was positioned as somewhat of a debate between wide area and Wi-Fi. I think a year ago the differences would have been a lot starker between what Sky was presenting and what I'm about to say. But the world has moved on, and like any new technology Wi-Fi is hitting its level. We're starting to see the places where it is adding value versus some of the places it was initially hyped.

The key learning for me, as I began to write my little travelogs, occurred when I started to hit different Web sites like Glenn Fleishman's Wi-Fi Net News, places like that, which from a wide-area world I'd never been exposed to. One of the comments I made is that if you have one individual and one TV station and all he can watch is PBS for news for years on end, and you have another individual right next door who gets one television station and all he can watch is Fox News for years on end, they're going to have very different worldviews as a result of that perspective. I think that some of the debate and dialogue that's going on is taking two rapidly evolving technologies and trying to reconcile where the world will end up.

Qualcomm believes in Wi-Fi. We utilize Wi-Fi all the time; we believe in Wi-Fi in the enterprise at home. We have spent over $300,000 deploying Wi-Fi in most meeting spaces in public areas at Qualcomm. I have Wi-Fi at home: I have six direct employees and they all have Wi-Fi in their homes. We utilize Wi-Fi as a technology on a daily basis.

The key point here, though, is to differentiate the café situation, where somebody just takes an inexpensive access point and puts it on top of maybe an

* Senior Vice President, Marketing, Qualcomm

existing back haul or on top of, as Sky mentioned, business-grade DSL. The costs of deploying a hotel such as this one or the costs of deploying an airport—there's some recent information on Glenn Fleishman's site about it—can be a multi-million dollar expenditure, depleting all of the infrastructure to enable those access points to be hung. Again, two very different business cases. If you go to Stay Online's Web site, they have some white papers showing what the costs are to deploy in a mid-size hotel. The reality of a large-scale deployment is very different from a café operator's.

Some of my writings were prompted by the fact that I was having a difficult time given the amount of scrutiny that wide-area data gets with the ubiquity of articles talking about going to your local coffee shop to get 11 megabits per second up to 300 feet from the coffee shop. The reality of it is that if it's a T1 connection, which is costly in the high hundreds of dollars per month, you're going to get a maximum of 1.5 megabits per second. If it's a DSL connection, you're probably going to get in the low to mid hundreds of kilobits. Again, it's not free. Coverage is limited.

My key thing here is the travel-to-compute model. I agree completely with Sky that places where people naturally compute are evident in places where Wi-Fi will thrive. However, a lot of the business models and announcements, be it a McDonald's or Borders, are predicated on the assumption that someone is willing to drive five minutes or walk five minutes to get to an access point to compute. I like taking the drive five-minute example. Does that include the time to your car, finding a parking spot, getting out of the car, walking to the access point? I think that's a triviality, but the point is that in this room I would assume we have near ubiquity of digital cell phones, and I would ask the folks in this room how willing they would be to walk five minutes or drive five minutes to make a phone call. It's our worldview that whether on Wi-Fi or WAN systems your desire to have Internet connectivity from a broad range of devices will need to be ubiquitous. A number of places like airports, hotels, etc., have multiple landlords and multiple costs that will drive the costs up as well. Billing and roaming are issues that Sky is obviously looking to address and moving forward on.

It's our worldview that either on Wi-Fi or WAN systems your desire to have Internet connectivity from a broad range of devices will need to be ubiquitous.

I recently asked several hundred IT managers at a CIO conference what their range was, and it was rare that anyone got over 100 feet in a real life 802.11 deployment. In our multiple hundreds of access points we use about 90 feet; πr^2 against that to determine the area of that access point is about 25,000 square feet of coverage per access point. A single suburban cell site of about 3 miles of coverage translates into 750-million square feet per cell site. So when you see the announcements of another ten fast-food restaurants being covered in Wi-Fi in a given city, what's not being told is that the Sprints and Verizons and the like are putting out thousands of cell sites each year just to fill coverage holes. Each one of those cell sites is roughly the equivalent of 25,000 access points. Based on Gartner's figures, using that ratio, by the end of 2006 the amount of aggregated coverage is the equivalent of four cell sites in the whole continental U.S. Using Sky's numbers of 300,000 hot

Public WiFi Service Limitations

- **Data speeds**
  - Limited by backhaul and multiple access scalability
  - 11 Mbps becomes irrelevant when connecting through a T1/E1 (~1.5 Mbps), DSL or cable modem (300 – 500 kbps)
- **"Hotspot" coverage**
  - Very limited
  - Predicated on "travel to compute" model
- **Backhaul costs**
- **Landlord fees/revenue sharing**
  - Perceptions of ultra-low service fees are incorrect
  - Hotel room phone example
  - CTIA IT show / T-Mobile example
- **Billing issues**
  - WiFi roaming is in its infancy, need for multiple subscriptions
- **Barriers to entry are few**
  - "Java Joes" can provide free access next door to a Starbucks/T-Mobile

TELECOSM 2003 August 24, 2003 QUALCOMM

Mainstream Users Expect Ubiquitous Coverage

- A single 802.11 access point covers roughly 25,000 square feet
  - One or more APs consists of a WLAN "hotspot"
- A single suburban 3G cellsite covers roughly 750,000,000 square feet

Number of public WiFi access points (est.)

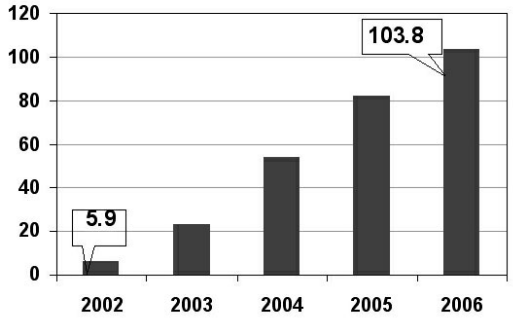

By 2006, estimated U.S. public WLAN access points will cover an area roughly equaling 3.5 cell sites

Source: Gartner Dataquest 2002,

spots aggregated around the globe, that's the equivalent of 10 or 12 cell sites of global coverage. Put that in perspective. Sprint alone has in the area of 20,000 cell sites in the United States.

Here's a prediction that I made that I think you'll see change. T-Mobile's own numbers show about 46 minutes of access per hot spot, all using a T1 connection. If any of you have businesses utilizing a T1, $600-$800 per month is not unheard of, or is sort of a midrange cost per month for those T1s in each of the Starbucks. In my mind you'll probably see that go quietly off their Web site at some point because it's just too expensive to provide, especially since they're averaging about one user per day. Again, those numbers will come up. There will be more utilization. The concept that you can have a thousand dollars or more in fixed costs for a hot spot remains challenging.

We agree with Sky that places where Wi-Fi will proliferate or places where people naturally compute are the airports and the hotels, but if you look on the directories of a lot of the hot spot folks, the access for those Wi-Fi hot spots are typically lobby access. Again, people will want ubiquity just as they want ubiquity on their cell phone service. There's a stat that I've been digging for that shows the majority of people traveling do their computing and e-mail at night in a less-than-dressed condition in their hotel room, sitting on the bed or at the desk. It gives you some very scary imagery if it's 10:30 at night and everyone's down in the lobby in their underwear on their access points. I'd encourage you to go beyond the announcements that have been made and actually check out some of the providers to see the degree of coverage that exists, which will improve of course. Bottom line: to get Wi-Fi access in the course of a business trip, you have to arrange that trip specifically to find places where there's access.

To get Wi-Fi access in the course of a business trip, you have to arrange that trip specifically to find places where there's access.

Verizon, on the other hand, and Sprint does the same thing, has a PC MCI card for CDMA2000 data, which provides two to three times real dial-up speed, not the 56K that none of us has ever gotten on our modem, so anywhere in the 50 to 90 kilobit per second range. It costs about $80 a month for unlimited access, better than dial-up speed. Again,

as Paul Jacobs said, he downloaded 40 megabytes at his house. It's not dollars per megabyte; it's ubiquitous, better than dial-up speed. Verizon has announced that it's rolling out the next step in the technology, and I apologize for the acronym—CDMA 2000 1x EvDO—which will be 2.4 megabit, peak, but realistically in the middle hundreds of kilobits per second in San Diego and DC in a broad metro area. They haven't announced pricing yet, but the system will be going live over broad areas of DC and San Diego. Watch that space because I think it will sharpen this debate since consumers will want the highest speed as ubiquitously and inexpensively as possible. Operators, as Sky has indicated, will look for ways of minimizing their cost per bit while providing the best customer service possible.

If you haven't tried wide-area data, haven't tried high-speed wireless, I'm encouraging you to do so.

In the long run, just as you get upset if you can't make a phone call anywhere, in our view that's the way you're going to want to access the Internet—as fast as you can in as many places as you can. Internet on your phone, Internet on your PDA, Internet on your laptop. In that respect, the operators will look toward ways of making the technologies coexist. If you haven't tried wide-area data, haven't tried high-speed wireless, I'm encouraging you to do so.

Spencer Reiss of *Wired* magazine

Cramming for the next session.

Jim Mullens just couldn't stay in his seat.

Qualcomm founder Klein Gilhousen steals the stage from Brian Modoff (Deutsche Bank), Tom Pollard (Texas Instruments), Omid Tahernia (Motorola), Jamie Cummins (QuickSilver), and Bill Colleran (Impinj) on the Communications Chips panel.

The Unwired Revolution

monday

august 25

12:30 pm

Paul Otellini:*

The world of headlines is reflecting some of the mania that's going on around Wi-Fi and wireless in general. It would not be current if I did not bring you a copy of that paragon of accuracy, the *San Francisco Chronicle*. There's an article in there called "Wi-Fi Hits the Spot." It's about this guy who has a coffee shop in North Beach who put in free Internet access, and his traffic has gone up by 30 percent. Sky told us a similar anecdote. What's going on here? What's happening in this revolution?

That's what I want to talk about today, the Unwired Revolution. It's not just about Wi-Fi, but the whole move to unwire yourself and make things more available.

I'll cover three things. First, why is Intel driving this notion of convergence of computing and communications? What's behind it? What's Moore's law's role in it? Second, Wi-Fi itself, as an environment and as a way we'll use computers. And, third, a little bit about the future. What's next? What's after 802.11?

Let me begin with the construct that we build our products around at Intel. The idea of convergence has probably been talked about for longer than I've been at Intel. Many companies and writers have had different notions of what convergence means. The generally accepted view, though, is that it's somehow integrating *communications* and *computing* together. NEC's corporate slogan picked up that "C&C" almost three decades ago. But it is critical for Intel, and critical for many of you in the room, because technology has caught up with the dream. In our business model, our vision, and our product planning, we have a very simple premise. We believe that all computers in order to be useful have to communicate

* Paul Otellini, COO, Intel Corporation

and increasingly all communication devices will have increasing degrees of intelligence to make them more useful—a very simple premise predicated upon what can be built on any given chip. The idea of radio-free Intel is not that far away. It is absolutely conceivable to put a radio, multi-band, selecting any network that's out there, on the corner of every microprocessor or microcontroller that we build. It's absolutely in the cards and in our product development plans. We're going to add communications to all computing devices and increasing degrees of computing to all communication devices.

Seven years ago Andy Grove made a prediction. This was pre-Internet. He talked about the notion of what a billion connected computers would do for us around the world, how it would change what we do with computers and how we live. Four years ago, Craig Barrett augmented that and said that it's not just about a billion connected computers; it's also about a billion connected handsets. These kinds of things are not just opportunities to sell silicon but, increasingly, opportunities to sell services and to deploy very interesting applications. If you look at what happened since Andy first made his prediction, we're there—a billion data-enabled phones. And PCs will get there in the next year or so. What's next?

If you extrapolate the trends and the technologies and the price points and the costs another six years or so, what does 2010 look like? It's highly likely we will see a billion and a half broadband-enabled PCs. Broadband networks are reaching all of us wherever we go. On the handset or PDA side—things you carry around in your hand—it's likely there will be 2.5 billion connected around the world. As you know, most of the growth in the handset markets is in emerging markets where cell phones are the first and only means of communication. Most of the growth in our businesses in those areas has been in China, India, Eastern Europe, and so forth. But the new notion is the kind of performance that we'll put inside those handsets six years from now.

The top-of-the-line handset chip from Intel today has the performance equivalent of about a 500-megahertz Pentium II. I suspect more than one of you in the audience still has one of those. That was our state-of-the-art product five or six years ago. By 2010, we will put the performance equivalent of a 4-gigahertz Pentium IV into the handsets. This is what Moore's law is all about. They will be single-chip, multi-radio, multi-protocol, very, very computive-intensive machines. It's what drives us. Our product planning cycle is three to five years. This is where we're headed. These are the kinds of products that we'll bring out for our customers, and many of you will have the opportunity to start thinking about the application space, the service space, to take advantage literally of billions and billions of machines of various types that will always be connected in a very digital fashion.

It is absolutely conceivable to put a radio, multi-band, selecting any network that's out there, on the corner of every microprocessor or microcontroller that we build.

The combination of wireless and all its manifestations and the combination of compute silicon across multiple architectures is what will enable convergence. It's that simple. One of the reasons it hasn't happened before now is that Moore's law didn't allow it to happen. It was just too expensive or too

complicated. Now, the networks are there, the costs are there, and you're seeing explosive growth, which is why wireless is so interesting. It gives all of us as users exactly what we want. People ask me if Wi-Fi or wireless is the next killer app. I think it's the next killer environment. Wireless computing in any manifestation doesn't change what you do with your computer. It changes how you do it and where you do it. It makes it much more accessible, and it drives you to have much more on-time real-time data in the home, in the office, on the road. These networks will work together and be increasingly cheap and increasingly more powerful, simply because they are always digital and they're *all ways* digital. Both aspects drive the technology.

Wi-Fi is in danger of being over-hyped. I completely agree with that, and to some extent Intel may be part of that in spending multiple hundreds of millions of dollars on our Centrino ad campaign this year. But it's not just about a campaign about a chip or a set of chips in a notebook. It's a campaign about use models. We spent almost as much time enabling hot spots, verifying that these machines work when they turn on, can sniff out networks and work seamlessly as we did in developing the chips. Our venture arm was incredibly active with many of the companies in this room, focusing on getting the business models and the deployment out there.

Intel brought a catalytic effect to the market. We didn't invent Wi-Fi. We weren't even the first to market. We were the first to market in an integrated fashion with a very high performance microprocessor and chipset, and we put significant money behind it. There's an old adage that if there's a parade the best thing to do is to get in front of it. In many ways, that's what we're doing with the initial versions of Centrino and some of our handset chips. The *Economist*, which is not prone to exaggeration except in the area of politics, is an indicator of the almost foaming-at-the-mouth comments written about the technology. This will settle down as the real business of wireless and Wi-Fi becomes apparent. But I think we have to ask ourselves, Why is there all this excitement? There are a number of reasons. This technology is relatively easy to use. It

There's an old adage that if there's a parade the best thing to do is to get in front of it. In many ways, that's what we're doing with the initial versions of Centrino and some of our handset chips.

took me less than five minutes to sniff the network, launch my VPN, and get online out here [Lake Tahoe]. My assistant was sitting next to me in the lobby doing his Voice-over-IP call in that same time frame. A year ago it probably would have taken me two network engineers and a lot of help to make that happen. It's getting to where real people can do this very easily.

It's also interesting to me that this is an unlicensed phenomenon. This technology for the most part is not regulated. There are no people you have to appeal to in order to put up hot spots, public or private. You just deploy it. As Sky pointed out, it's really cheap, and the investment to get online or to broadcast online is very, very low. As a result, we're seeing organic growth like nothing seen before in the regulated environment in a long time. Guess what? It's broadband. Someone earlier today said that you will not be limited by the speed or lack of speed on your Wi-Fi network. Absolutely. That's so true. There are some elements in the home, though, that will consume what we're currently building, and I'll talk more about that. It's a standard and it represents convergence.

Public hot spots are in the 25,000 to 35,000 range this year, probably doubling next year. That curve, however, is nonsensical if you look at the aggregate of all the vendors selling access points today. There is a couple of orders of magnitude difference in the numbers, and the world is purchasing something north of 15,000 wireless access points per day. Per day. It doesn't necessarily mean public hot spots, but the deployment of the technology is moving much, much faster than the projectors are able to internalize what's happening around them—back to organic growth.

The catalytic effect of Intel jumping into this market, trying to get in front of this parade, is phenomenal. Dell is already saying that it will be shipping 100 percent of their notebooks by the end of this year with Wi-Fi. That's true. You'd be crazy not to buy a notebook with Wi-Fi. The comment earlier on the battery life of PDAs or handsets in an 802.11 environment was right-on. Anyone who has used one knows that it sucks battery life like there's no tomorrow. Intel will solve this problem. We have development underway to deal with this problem. A number of companies in the industry are moving to greatly accelerate how fast we take the power consumption down on these networks.

This technology for the most part is not regulated. There are no people you have to appeal to in order to put up hot spots, public or private; you just deploy it.

Where can you connect today? There are all kinds of stories out there. There's a city in Spain called Zamora, a beautiful, stone, medieval city, and rather than tear its streets and buildings up to install modern telecommunications, it Wi-Fied the entire city. You can do Wi-Fi in the base camp at Mount Everest. Should you want to, you can Wi-Fi a McDonald's, Times Square, and elsewhere. Most Starbucks. Eight thousand cafés in Seoul, Korea. I have a house at Lake Tahoe, and I have a line of sight that comes across the lake 13 miles and Wi-Fi inside the house. It's here and it's easy. It has a few issues though: security; how easy is it to connect; the roaming agreement, which I think is in the critical path right now for massive wide-scale deployment; spectrum availability; power consumption; and where the heck are all

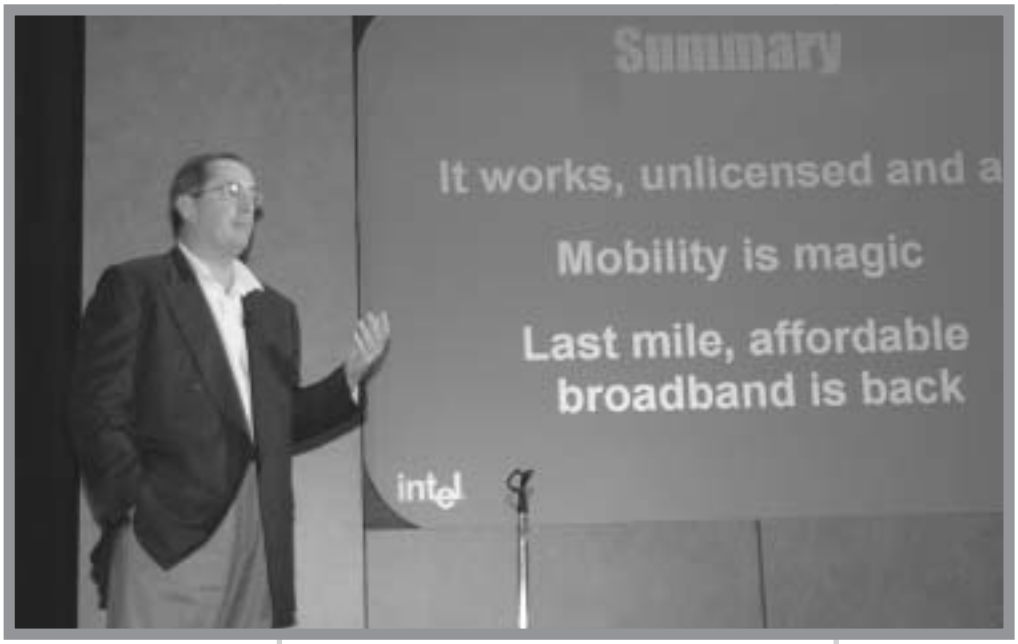

these hot spots. But there are solutions, either existing or in the works, for every one of those problems.

Security has made real, real progress. Wi-Fi protected access (WPA) is in widespread use today. There's a new variant of the 802.11 spec that is called 802.11i, which is security enabled that will be sanctioned this year, and you'll see that deployed into existing networks. It's simply a software upgrade to the 802.11b, g, or any network out there. Subscribers are growing. Just go into a Red Carpet Club anywhere in the country. Thirty of this country's major airports are already lit up, and it's happening worldwide. With the wireless radio on, we're getting three to four hours of life out of top-of-the-line Centrino notebooks today. Those of you who have used Wi-Fi in the past know that you couldn't do that before unless you found a plug somewhere. What we need most, in places like this ballroom, are outlets. Many of you would like to be connected but probably don't want to be because the wireless radio inside that notebook is going to take your battery life right down. We have to address this and we are.

More spectrums are available. The World Radio Conference just convened in Geneva and allocated an additional 455 megahertz to the spectrum of the 5-gig band. You heard from a number of speakers today that costs are coming down. T-Mobile and Boingo are under $30; if you're a broadband subscriber of Verizon, it's free. It's cheaper by the day and getting easier and easier to use. People are using it everywhere: 20-million estimated users in the enterprise space, estimated to go to 120 million by the end of 2006. At Intel we are not building any new buildings that don't have wireless built into them. It's the cheapest way to network our buildings going forward. I've got to believe that everybody else sees the same economics we see. We use it today in our factories. We have Wi-Fi-enabled PDAs in one of our most advanced fabs, and it has allowed our technicians to improve communications. The data entry accuracy has gone up because you don't have to go from the machine back to the workstation and record the data. There's a transcription potential error rate when that happens. Now you can do it right there from the spot. We see significant increases in productivity. When you look at most of the studies that talk about what the ROI for

There's a city in Spain called Zamora, a beautiful, stone, medieval city, and rather than tear its streets and buildings up to install modern telecommunications, it Wi-Fied the entire city.

cap-ex is for buying wireless-enabled notebooks, they tend to say that if you get 30 more minutes a day from your employees, it's worth it. We're seeing substantially higher ROIs than that by bringing this technology inside our factories.

I think it's also interesting to note that this is not just a handset or a client or a PC phenomenon. Let me make a proposition. As the Internet was to the growth and the surge in demand for edge servers, I think that wireless has a similar potential to drive demand for database, something I'll call OCC, or occasionally connected computing applications. The applications I mentioned in the panel earlier are mobile-aware from the get-go. Let me give you an example of what wireless data services can do for servers: DoCoMo. DoCoMo, as you know, in Japan is the i-mode phone: 75,000 data transmissions per second; 800 million transmissions a day; 20 transmissions on average per user. To run all that, they have 400 servers at their i-mode center, supporting 400 terabytes of information and 50-million simultaneous users. This kind of service drives huge demands on databases, on applications, and on the servers that will serve those up on a worldwide basis. Ten million users happening today in the home. Most of those wireless access points I showed you earlier are going into home or small business environments—expected to be 90 million by the end of 2006.

What's happening with the other devices in the home? You may have noticed that there was an announcement made a couple of months ago regarding the DHWG (digital home working group) Intel led; it's a coalition of computer manufacturers, consumer electronics manufacturers, and software enablers. Think of getting Sony, Microsoft, Intel, and Dell to agree on standards for interoperability inside the home. The person from Intel who led the campaign received the Nobel Prize. It's a tremendous leap forward. What it means, quite simply, is that as new CE equipment is built, increasingly it will be made for a wireless network-enabled home. This is where some of the bandwidth issues become more troubling. While it's certainly true that with 802.11g you can move pictures around and maybe even do some low-quality video, but if you want to get high-definition signals moving around the home to multiple devices, you need to increase bandwidth. That's one of the issues that DHWG is working on. The limit, though, in both the enterprise and the home, will be the back haul: the deployment of broadband to the premises. In the U.S. this is being driven principally by the cable folks and companies like Verizon. But if you're outside the U.S., and in many cases even inside the U.S., broadband connections are very difficult to find. Try getting a broadband connection at your house at Lake Tahoe. I tried. You can't get DSL up here. We had to solve the problem by beaming something across the lake and creating a mesh network for our neighbors. Try getting it out in the middle of Iowa. Try getting it in the valleys of south San Jose. It doesn't exist. You have to deal with this issue in the most cost-effective, rapid fashion, and I think in the near future it's going to be wireless.

Think of getting Sony, Microsoft, Intel, and Dell to agree on standards for interoperability inside the home.

There's a new technology that is also in danger of being overhyped, but it's one that is incredibly interesting, WiMax. It's 802.16; 150 times the speed of DSL; up to 6 miles; it has a 30-kilo-

meter range, potentially in an unlicensed spectrum and substantially less than that in a licensed spectrum. What we're seeing is that people are very interested in deploying rural networks. In many ways, it's described as a metro area network; it's really a rural area network in terms of where the deployment will be. It's getting people access who don't have it today. The economic model is relatively cheap: you can stick these antennas up on the same kind of cell site antennas that are out there; feed them at the same backbone; and deploy your broadband very, very cost effectively. It's not perfect and not yet ready for prime time. The spec will be solidified early next year, and you'll see product start rolling out. Now it's stationary: it's 802.16a, but there's a variant of it that's called 802.16e, which adds mobility. I've talked to people in Europe who operate telecom networks in countries where they use trains a lot, and this one telecom network that already has cell site coverage along the tracks wants to put 802.16e antennas on its existing cell site towers. As the train goes down the track, you can have a seamless connectivity: 16 to the train; 802.11 inside the train. Very cheap, very pervasive. And you can do it in a car and you can do it for telemetrics and for telematics. This particular technology operates in licensed spectrum in metro areas and in unlicensed spectrum in rural areas, which makes it very, very exciting. It's line of sight and non-line of sight access as well.

The technology of the unwired revolution is terribly interesting because it works, it's unlicensed, and its mobility is absolutely magic.

Wrapping up, the technology of the unwired revolution is terribly interesting because it works, it's unlicensed, and its mobility is absolutely magic. Once you've gone Wi-Fi in your home, you'll never go back. Once you've gone broadband, you'll never go back. This is

a huge change for all of us who have lived in the world of wired computing for so long. Probably most interesting is the economic tension that this technology puts on the business models of the incumbents. The technology will deploy. It will all coexist. The cost will come down. And the cost will come down dramatically as data-intensive applications go up. That really is what most of us are interested in, in terms of the future of the business.

George Gilder:

Sky Dayton said the revenues from Wi-Fi chips are going down because of their drop in price. How is that affecting Intel? Has it gone down at Intel as well?

Paul Otellini:

The guy who runs our business said that 802.11 was the first technology to become commoditized before it was introduced, and that's so true. It's one of the reasons why we talked about the radio-free Intel notion. Putting radios on these chips will be similar to putting graphics in chipsets or floating point engines in microprocessors. Ultimately Moore's law subsumes all those functions. In the near term, there is tremendous pricing pressure on Wi-Fi chips, but that's life. Welcome to the semiconductor industry.

Bret Swanson:*

Could you give us a little more detail on your efforts to put more products into handsets and other mobile devices as opposed to the flash memory, which you had been doing?

Paul Otellini:

There's an initiative at Intel we've had for about four years now called PCA (personal client architecture). What PCA is all about is turning handsets and PDAs into re-programmable machines much like you've seen in the PC model. Let's replay the PC movie along the same construct. Up to this point in time, pre-PCA, every new handset, every new PDA had custom ASIC silicon spun, which meant you had to change the software and so forth. We propose to do a Moore's law kind of vector of performance on a common architecture where we can have multiple

Putting radios on these chips will be similar to putting graphics in chip sets or floating point engines in microprocessors.

** Gilder Technology Report*

generations and embed that into increasingly smart phones and PDAs. As a developer, that's great. You save money building your own device. As a service provider, it's great. You now have a common software platform. What we're doing differently than others is to develop a number of tools for developers to take their applications from the PC space, where there is quite a bit of work going on, and port them in a matter of days over to PDAs and handsets that use this architecture. You can re-purpose all the application work, and the tools know how to scale it down so that it knows it only has a three-inch screen, or knows it only has a certain kind of connectivity, and so forth. It's more than a silicon opportunity for us. It's a silicon and services opportunity, a software development service, that allows us a to have a much bigger footprint in the 400 to 500 million handsets that are out there per year.

When I talked to carriers in the last year to year and a half, I found they had gone from "heck no, this is death" on Wi-Fi to embracing it.

Brian Modoff:*

Question on WiMax. We've looked at the protocol. It's got multiple physical layers to it: CDMA; spread spectrum; RFDM; an alphabet soup of protocols. Aren't standards supposed to be standards where there's a set of defined parameters that drive it? Another question: you're missing some other large players in terms of supporting the protocol, players like Motorola with their canopy. They're saying we'll go with what we have and we'll see how that goes. Can you talk about where you see this standard going? Do you see it as finally making a decision and choosing access technologies?

Paul Otellini:

I think there are two parts to that: the standard and the business model. We don't like the standards. We were on the committees with everybody else, and as you know, the ITU standard-setting process is convoluted at best. There are lots of people who have axes to grind. The discussion of the next standard, 802.20, is a reflection of that because all the 4G arguments are coming to bear. God knows when or if that will ever get out as a result. Having multiple protocols and complicated requirements is actually good for our business model. We can throw compute power at it. We can throw transistors at it. I actually don't mind that particular problem. It gives us an advantage that others may not have in the silicon space. It

Not so sure Gilhousen is buying all this Wi-Fi mumbo-jumbo

* Deutsche Bank

will debut itself and work itself out; I'm quite convinced of that.

The business model side is interesting. When I talked to carriers in the last year to year and a half, I found they had gone from "heck no, this is death" on Wi-Fi to actually embracing it. They are very interested in 16, particularly 16e because it augments their existing infrastructure and business models. Back to the Qualcomm discussion earlier, in many cases companies like Qualcomm, or even Motorola, end up having to build handsets that their customers demand. If their customers are going to deploy these technologies because it's good for their business models, they end up having to support them.

Andy Kessler:*

I followed your company as an analyst for years. In the past, Intel sold to big, ugly, slow companies like IBM, and I know you almost personally put Michael Dell in business and Gateway in business and these alternate PC companies that ended up really growing the business. Do you foresee the same thing happening in this space where there's a bunch of big, dumb, slow carriers, and you're going to help invent a whole new set of carriers to sell to?

Paul Otellini:

Well certainly on the handset side that phenomenon is happening. The earliest ones designed were PCA, the architecture I explained in an earlier question, and came from second- and third-tier carriers, principally in emerging markets like China. But guess what? China is now the largest market for handsets. So companies that want to displace a Motorola or a Nokia are absolutely embracing this, and the carriers are adapting to the model around them. As I said, I've seen a marked shift in my discussions with the existing carriers in the last eighteen months. Given the deregulated environment, particularly offshore, starting new carriers has strong possibilities. One of the key assets the big carriers have is their customer base. As long as they don't preclude their customers from using these technologies by making it difficult with the roaming agreements, for example, I think it's going to work. I believe the dam was broken last month. There was a very small article about AT&T Wireless and Sprint agreeing to do roaming for Wi-Fi and sharing their customers on a common bill. Once that begins to happen, everyone's back in the game again.

Given the deregulated environment, particularly offshore, starting new carriers has strong possibilities.

Unidentified (ATTENDEE):

You talk about a six-mile distance with WiMax and you also mention line of sight, but that would not be a six-mile radius, if I understand you correctly? Also, would your chip that will have all-radio technologies in it, including CDMA, require a license from Qualcomm?

Paul Otellini:

Answering the second question first, yes, it would, and at some point we'll have to deal with that. Right now, as I said earlier, most of our business and early design wins tend to be in the GSM world and its progeny, because it's going on now in India and China and

* Former hedge fund manager • Author, *Wall Street Meat* and *Running Money*

Judy Canfield

places where there are other standards. In answer to your first question about 802.16, as I said, there's a two-by-two matrix. There are licensed and unlicensed elements of it, and there's metro and rural so that you get a line of sight and a non-line of sight environment. It will have a much wider carrying capacity in the unlicensed rural spectrums than it does in the metro spectrums. It's that simple.

Judy Canfield (ATTENDEE):*

I'm wondering what the difference is between a Pentium IV and the Centrino chip in an already Wi-Fi-enabled laptop computer.

Paul Otellini:

Sure. The processors are different; the chipset is different; the radios are different, and apart from that they're all the same. The Pentium M processor, which is the processor in Centrino, is of higher performance and lower power than any of the Pentium IV mobile processors we build. It has longer battery life and all that other good stuff. The chipset is more highly integrated and power-aware. Wi-Fi chips are also power-aware, so we use the microprocessor and the software to manage the power across all of the chips, which gives you longer battery life. That's how you get the three to four hours with the radio-turned-on phenomenon. In addition, it's just faster for all the applications and benchmarks.

Judy Canfield:

So it's faster doing regular processing?

Paul Otellini:

Yes, if you turn off the radio, it's still a faster notebook.

* Director, Envision

Adam and David Eslinger contemplate adopting a less-fortunate, fatherless attendee.

Kessler travels with his own mic for the Karaoke session.

The panel ended 15 minutes ago and they're still taking notes.

Telecosm inspires thirst.

The Blunt Edge of American Broadband

the future of the broadband local loop

Douglas Bulleit
Chief Strategist, BellSouth

monday

august 25

1:30 pm

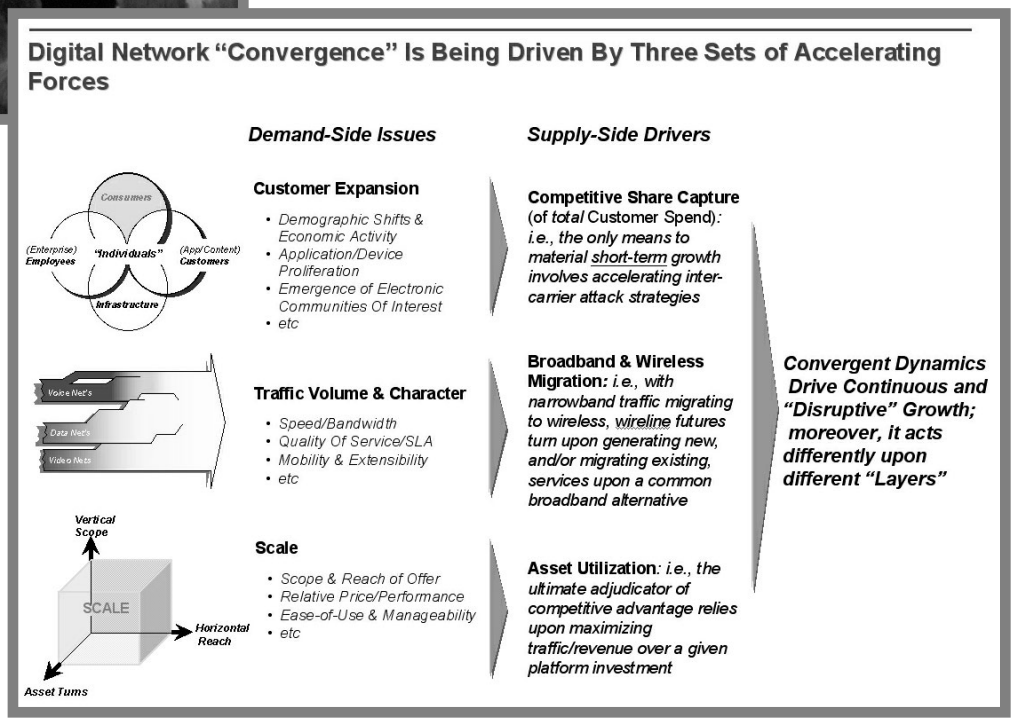

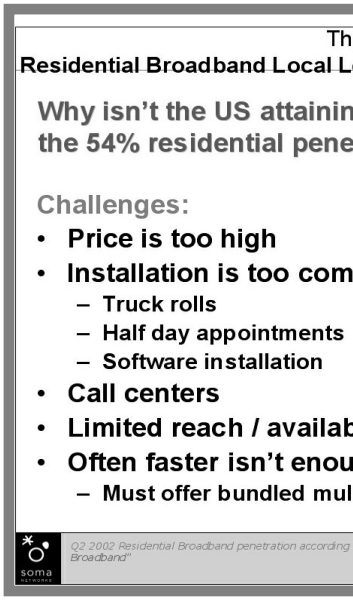

Yatish Pathak
CEO, Soma Networks

Communications Chips

can wireless chips tame the protocol zoo?

monday

august 25

2:45 pm

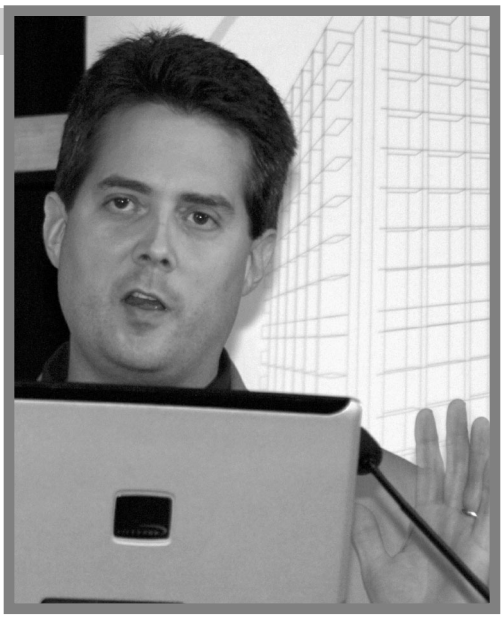

Tom Pollard
Worldwide Marketing Director
Texas Instruments

OMAP
TEXAS INSTRUMENTS TECHNOLOGY

Integration at the Chipset Level

Increased functionality, smaller size, reduced power consumption

Why Integrate?

- Reduced cost-per-function
- Cost reductions at the system level
- Increased performance for multimedia features
- Longer battery life – better talk and stand-by times
- Smaller form factor to enable small, sleek mobile devices
- Easier design and test for wireless designers due to new turn-key, ready-to-manufacture solutions

- ***Coming Soon – the Single Chip Cell Phone!***

2.5G/3G Terminal Functions

Texas Instruments

Infra-destructuring and Awareness

when all things are on the 'net'

monday
august 25
3:45 pm

Greg Papadopoulos:*

This is going to be a really different talk from what you've heard so far. I have four basic concepts to introduce to you.

My interests and a lot of what I've been thinking about, in particular over the past year since the last Telecosm, are really the implications of what happens when we get beyond what some of these discussions have been about, of those devices that connect people to the network, and into what will happen when we get to the point of simply connecting things to the network. Certainly RFID (radio frequency identification) is a good leading example of this kind of technology. But a lot of the technologies building ever and ever cheaper radios are going to seep into all products, no matter what.

Since Moore's law is the favorite exponential for our biz, one of the interesting things for me is to track one of the inverse consequences: what does it cost you to connect something to the network? This cost is the cost of engineering the device to do it, and, of course, that should follow a similar exponential decline. From a decade ago it cost maybe fifty or a hundred dollars to get an NIC card to plug into your PC or your server; today it costs about a dollar to get something hooked up on the network. A decade from now I fully expect it to be a few pennies.

I saw something from Cypress last week. They have a wireless USB that has a 1,000-foot range for maybe a dollar or two. I believe that the future of everything connected to the network is going to be wireless, simply because the connectors are going to cost too much, or wiring will cost too much. And we all

* Executive Vice President & CTO, Sun Microsystems

know about RF tags. I think of these as an excellent example of something on that curve that is giving network presence to atoms. You paste this stuff on atoms, and I then know something about them. I may know a unique ID and the geographic coordinates, because I know the coordinates of the reader that's reading them, but they'll have more and more information over time. RF tags are an interesting lens to what it costs to give a signature to something on the network. Maybe it costs about a quarter today. Looking at the revolution from a business point of view, we've explored in detail what the value is of connecting, mailing, and transferring files to community building, which I think is really the essence of all these propositions built around people. You build various flavors of communities, and if you're a business you're building communities of your customers as well as of your suppliers, your employees, shareholders, etc. The leading concept in my mind is that when you start to connect other things everything that we bring to the network is about creating awareness. It's the concept of taking automation, which is a core of IT, and moving it outside the walls of the enterprise into the life cycle of something that someone made—manufacturers and others—and becoming aware of what's beyond the walls of the enterprise.

Some of the four concepts I promised are strangely named, but I hope you'll at least be amused or have some new words you can abuse and throw around.

Awareness. This idea that network entropy, which is a destructuring effect in which the network has to tear things apart in pieces, is an important concept. Think of bit mass of things that have atoms in them that will increasingly have bit signatures. Think of it as RF tags to begin with and then how this can remake infrastructure. It's the frontier of automation.

There's a very interesting debate going on now. Certainly, the *Harvard Business Review* article a couple months ago fueled that debate. Is IT in a post-innovation state? Is it now all about driving the cost out of things? And the commodity argument: we've done everything that we need to do; everything is just a commodity; we need to drive cost out. I have a real pet peeve with the word "commodity" because, first of all, it doesn't apply in the sense that we really mean commodity. And if you mean it in driving the cost out of computing, computing is always being commoditized. You better be able to do today what you did three years ago for a quarter of the cost. If you can't, you're not on the curve.

You better be able to do today what you did three years ago for a quarter of the cost.

If you think about this debate, if you look inside an enterprise at the different functions of HR, Finance, Sales, Service, Operations, R&D, etc., you'll discover that we have successfully automated them. In fact, in the last decade a lot of Sun's growth was a result of our being able to handle larger and larger scale enterprises until finally we could snap them in, run a single image of SAP, and run an entire enterprise globally. Well, we've done that. So now you better just drive the cost out of it. One of the things you can ask is whether this is a source of growth. Ask a CFO, "What would you do with a million times more computing capacity than you have today?" And what will the CFO say to you? "Close the books every millisec-

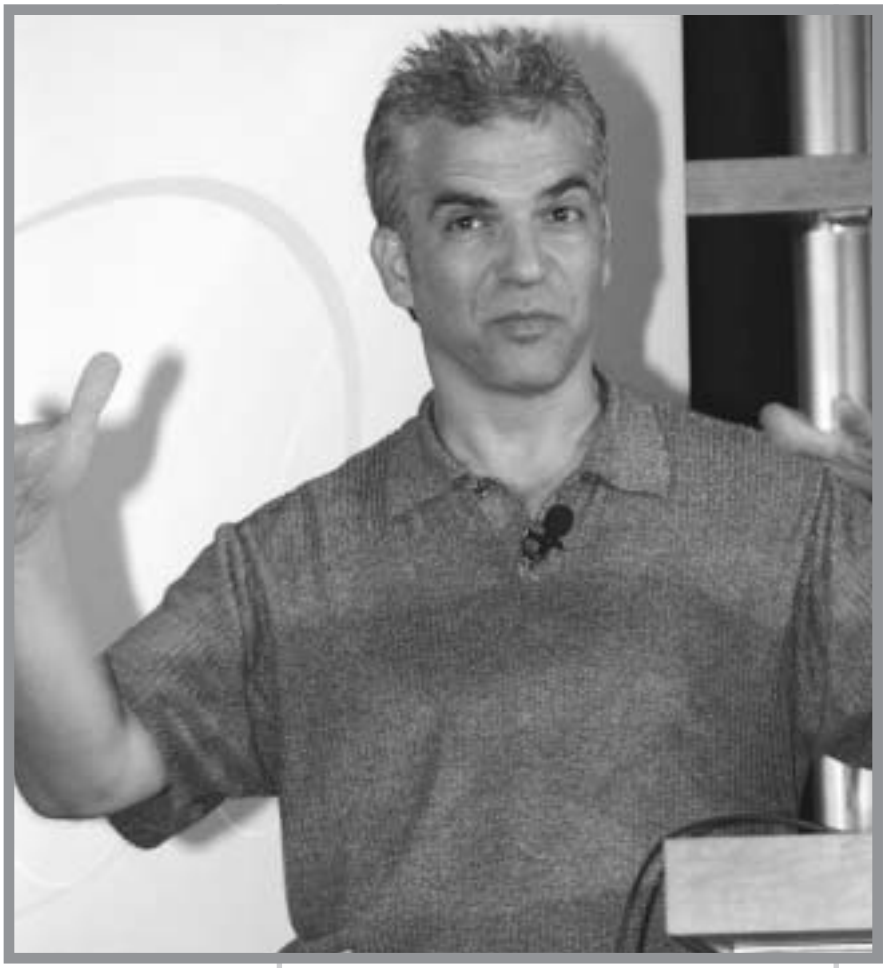

than you have today, he'll say, "Bring it on." In fact, he'll say a billion.

Quietly this reinvention of high-performance computing is taking place, and there will be a tremendous reawakening of the market in that space. But the other thing that's happening is this whole community space. It's essentially the infrastructure that supplies all these great radios and the wireless connectivity we're going to see.

But what happens when you connect everything? By everything I mean anything from really high capital good items like GE's connected aircraft engines to pumps and water purification systems, to refrigerators and washing machines, to the coke dispensers inside McDonald's, to packages with RF tags. What does that world really look like? It's this area where I can attach networked intelligence to things I've made as an enterprise. I get out of the mode that exists for most enterprises—99 percent of them—that once the product leaves the shipping dock, they don't know anything about it. At best, it's a replenishment order that may come in, or a customer sends in a warranty card, or someone goes out to a service center for installation. But there isn't a continuous relationship with these items. Think about the entire services organization. A customer calls and says, "This is broken." "What's broken?" "I don't know, it doesn't work." They roll out a truck so someone can check out the problem. You know the story. You think people have awareness—you think FedEx knows where your pack-

Ask a CFO, "What would you do with a million times more computing capacity than you have today? "He'll probably say, "Why don't you cut the costs of my computing by a factor of a million and be done with it?"

ond?" No. He'll probably say, "Why don't you cut the costs of my computing by a factor of a million and be done with it?" So what do you have to do in this business? You have to ask where the growth is going to be. If you're going to get anything out of this it is that if there's innovation and growth to be had, which will continue to fuel the capital cycle, where is that going to come from? There is an organization on that curve, meaning that if you walked up to someone in the organization and said that I will give you a million times more cycles

age is—and you know if you've ever had one lost that they don't. They have no idea where your package is. They only know the last time they scanned it. But ask the question: Where is my package? You will not get an answer of knowing where it is.

There should be awareness all the way down to the razor blades having RF tags. How do you imagine that going forward? Well that has added some very interesting flows to the entire network. There's the telemetry of the mass of data that comes up. It's the network turning around. It's not about distributing media out to people at the endpoints. It's about collecting understanding of what's happening in the world. Where is something? What is its state? Is that ink in the inkjet cartridge low enough that I ought to replenish it automatically? Does it still work? Then closing that loop and getting control back.

Of course, a connection that's interesting to a lot of people at this conference is that it's going to be intermediated by these global wireless networks, a task replete with all kinds of interesting problems. You not only have to drive the cost out of getting these things connected, but also the cost of a variable infrastructure you can afford to hook things to. And from my point of view, what are the pressures on the infrastructure and what kind of kit do you need to make all this stuff actually solve problems? And, of course, what are the programming models? I will dive into some pieces of this later, but now I want to go to my second concept, which may seem a little more abstract at first. In my mind, it is an incredibly useful observation that helps explain a lot of things that you're seeing happen in systems and maybe some unexplained things that look like they may be just transient phenomena.

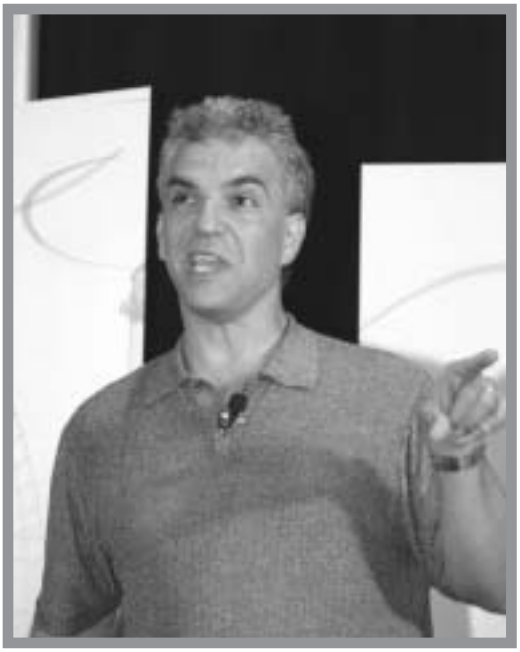

Networks want to decompose things that are connected to them, distribute their components, specialize those components, and then have a logical reintegration as you scale them up.

The second concept is network entropy. I named it after Rob Gingell, who's the chief engineer at Sun and a Sun fellow, who made this decades-long observation. It basically says that when you hook something up to a network, that network will inevitably erode the structure of the thing you hooked up. Networks want to decompose things that are connected to them, distribute their components, specialize those components, and then have a logical reintegration as you scale them up. I know that sounds very abstract. The cartoon image is if you take a monolithic system like a computer and hook it up to the network, over the decades the pressure will be to want to pull this thing across the network into various pieces. A simple example of that is computers not hooked to networks that have their local storage on them. Historically that represents a lot of computers, but it also

represents your laptop that is mostly not connected. What is the pressure when you hook a network onto it? Well you tend to re-factor the system. You decompose it; you say I'm going to separate off the idea of the client from the file system or the consumer of the data from the actual storage of the data, and you intermediate that with an IP network so that you then have a network file server. You've done the decomposition; it's been distributed over the Net. Now you can specialize. There are whole companies that have grown on that side. I made this point last year. Think about how you look at a Web model, and you see how it has pulled things apart to the point where now you've even taken the user interface, put that down into a browser, maybe carried it around on a mobile device, and then started to pull the pieces apart to application Web server tiers, caching databases, etc. It decomposes, but the question is why does that happen? Why would networks have that effect and what's the value? The answer is Metcalfe's law—there is a value of sharing the storage that you get as soon as you network. That's essentially what's causing these things to take place.

One of the big questions is what's the operating system? Do you want to engineer it from first principles, or do you throw people at the problem and have them all be custom?

One thing happening at a macro level is that things we used to call individual elements now have to be re-integrated. How do I take all this mass of stuff that I distributed out and understand it and manage it at scale? This is something we're struggling with right now in the business, and it's a big dividing plane in the way people think about it. You can think of servers connected to networks versus the idea of disaggregating all these functions, decomposing them, building a data center out of the server and storage elements, and wanting to distribute services out to wireless-connected networks.

How do you re-integrate? One of the big questions is what's the operating system? Do you want to engineer it from first principles, or do you throw people at the problem and have them all be custom? There's all kinds of syncopation that takes place with decomposition and re-integration.

This certainly is happening: networks will decompose the personal computer. PCs haven't changed their architecture. What is the decomposition? You pull the display out, which becomes a network display; there's storage referenced in the network; you get access to services; you have different modalities of getting into it. You can see this with network displays already. There's something we do at Sun called "Sun Ray," and there's the Windows terminal server out of Microsoft and "GoToMy PC." You can have quite a good experience of putting

a display, a compressed frame buffer, across the network. Interesting stuff and it leads to enabling the technologies to do this decomposition. It's all about state. For the longest time we got distracted thinking that the network equivalent of the personal computer was the network computer. And it's not. These things tend to be state attracters: I want all my stuff in one place. Why do you carry around your laptop? Because it's a state repository for you. You carry it around with you because you don't have good networking. If you did, you'd probably rely on the network to store more of your things. I carried this laptop here because it has a presentation on it. At some point in the future I'm going to come here without it because I know I can get access to it more reliably otherwise.

The equivalent of a personal computer in the network age is not a network computer; it's a personal network. It's the network representing your state. I don't mean a personal body network. I mean the Internet as a whole, having the overlay on it that no matter where I gain access to this network I'm able to get to my stuff. Address books don't belong on your phones; they belong on the network. They may be cached on your phone. I take photos. Where did those photos go? Where do they live?

There is this phenomenon taking place now, which is a kind of über device. I picked on Sony-Ericsson for this one. It takes the phone and the PDA and the camera and video game player and probably an MP3 player and whatever you have and mashes it into this sort of Swiss Army knife. And the Swiss Army knife is in fact the tool of last resort, when you need something especially to be done. Why has this happened? Because these things are really lousily networked to one another. If I have my address book sitting in the PDA and I need that phone number on my phone, it's the human Net that has typed those things in. If they were well networked, I wouldn't be compelled to mash them together to make it write a piece of software that would link the two together in the device. I really believe that. I think that as this device gets better and better and continuously networked, then networks will want to pull it apart. This is a transient time.

The equivalent of a personal computer in the network age is not a network computer; it's a personal network.

Hands-free devices are one of my pet peeves. Maybe Bluetooth will ultimately let me get there, but I only want to purchase one of those devices and have it understand what radio I'm near, what teleconferencing system will work, or have my own personal microphone. I shouldn't have had to be wired up to do my presentation. It's me and I want to introduce my voice to the conference. How does that happen? It all becomes fashion. I'm not too certain where that all ends up, but that's how cell phones, digital watches (remember them?), and PCs started, and they have all yielded to fashion. If you see a cool watch or what Steve Jobs has done or Nokia's urban phone, you'll realize that fashion is ultimately going to drive the consumer behavior in the spaces.

Concept No. 3. I'll try to bring you back now after this exploration of trying to help you understand these devices. What is the bigger future view of bits and atoms coming together and what does it lead to? It's a word I'm trying to coin, and you're being subjected to it, but it's this idea that things have "bit mass," this intertwining of bits and atoms. When I attach an RF tag to a device, to a box of

Cheerios, I've given that box of Cheerios some bit mass. It not only has its mass; it also has bits that go with it. There's a lot of physics here, but I just want you to think for a minute about this world of bits and atoms, with computing in the middle; in fact, communication and computing are the same thing. It's all communication.

Computer science talks about the top; physics talks about the bottom and all the circuits and devices and materials that we've heard a lot about. If you want to venture into this space and start to get very abstract about what happens here, you can go to all kinds of thinking that's happened over the decades. Landauer at IBM has done some basic work in this area. Speed of light matters because bits have location. There is a cost of computing. Ultimately we must be regulated about what it costs us to perform a computation. Is that computation reversible or not? Thirty years from now the biggest question we're going to have in our businesses will be the ecological one about how much energy we're expending to do all this stuff. It should have happened earlier. We should be collectively embarrassed about how much of that energy we've sucked out of the power grid.

When I attach an RF tag to a device, to a box of Cheerios, I've given that box of Cheerios some bit mass.

Bits do have location and the atoms supporting them have location, which causes all kinds of interesting things to happen. I want you to think about these bit mass examples. You can think of things like storage, smart cards, and memory sticks as examples, but anything that has any material that we make, any device, is going to have bit mass and connection. If you have at some point any digital component, anything that you build, it will cost you more to take it off the network than to let it find the network. All the circuits and everything you have will be around just wanting to do that, and you'd have to turn them off. Let me give you a more topical example, and that's the bottom of the space shuttle [Columbia], a missing tile. Imagine every tile has bit mass. They are dimly aware of what they are and who their neighbors are and what they've been through. If you have the misfortune of finding one of those tiles on the ground, you could ask it a question such as, "What happened to you?" I think that will be the nature of objects. You might want to talk to your house after an earthquake and ask it how it feels (something I worry about).

I actually wonder whether bit mass density follows network flux just as biomass follows solar flux. If you were to plot out the density of wireless bit densi-

ty over the years, I think you'd find the things that contain bit mass would follow that as well.

My last concept is a real tease, and it could destroy at least a year of your life just thinking about it. What happens when I combine the two concepts of things imbued with digital state and their becoming networked? Is it wise to take things we call infrastructure and tear them apart, separating the control from the function and doing that as a dynamic bonding? Sounds mysterious, doesn't it?

Here's my problem of the day. My sprinkler system at home has half a dozen different controllers of different pedigrees. Have you ever wandered down to a particular head and said, "The petunias are getting too much water and the geraniums are not getting enough." You can't control each head individually. When that system has been laid out with the controller, the valves, and the piping in the ground, there is a binding of the structure of the network, basically the watering network, with my control of it. All those sprinklers go on and off, and they're going to be controlled the same amount and the same number of minutes. If you could control each of the heads separately and had them networked, you could take that apart. You could separate them off and have a different kind of control. If you imagine this networked sprinkler head where you have an individual sprinkler head that has an IP-network connection, maybe a moisture sensor locally, some valves so each one can be turned on and off independently—then you would redesign the sprinkler network. Basically, you would build a water plane, and every sprinkler head would get plumbed up. In fact, it could be redundant. The valve is clever in that it can isolate broken pipe sections. Then you would have a network

and a way of teaching each head what its function should be. You've essentially taken that structure apart; you've destructured it. You can ask an interesting question: how do you identify each of those heads? That's one of the key concerns: there are all kinds of security issues. You don't want your neighbor's kid hacking your lawn. It actually turns out that they're not IP addresses. The sprinkler head needs to authenticate its ownership relationship to you and it can't

The last concept I'll discuss is a real tease, and I can destroy at least a year of your life thinking about it.

log into something. There are people currently working on this issue and doing some interesting things.

Here's where I'll give you another year of thinking about how this might affect things. Think about the basic construction of buildings, something Neil Gershenfeld has been doing with Internet Zero. Much of the costs of wiring a building entail doing a wiring plan and wiring each switch to the appropriate fixture in each room that's laid out. Instead, if electricians just came in and wired everything together and then returned later and did the teaching—that switch there is about these lights up here—if you kept that binding you would reduce a lot of the construction costs. The Library of Computer Science at MIT has RF tags on all its books, meaning you don't have to reshelf books in Library of Congress order. Why do that? The Library of Congress system was just an index for finding them. And I'm sure you can save money on the way that you organize things. You won't have the problems of lost books, but more interesting, you can have a librarian say, "Here's a set of cultural topics right now, such as governance in California, and here's a set of books related to that." They aren't going to come from the same part of the shelf.

Ultimately this line between bits and atoms is going to blur.

Think of gas stations, certainly a lot of retail distribution, the whole postal service. The post office should be a router. If I want someone to ship something to me, I could simply hand out a random tag to him and the post office would look up where I am and rout it to me. You can go on and on.

Let me conclude here. As I said, I don't have any answers. Maybe I twisted you a bit into thinking about things differently. It's inevitable that we're going to be connecting. You can't help but connect things onto the network. The idea of business awareness means extending business processes to things that have been manual, which we haven't been able to afford to do before, extending those out to the servicing of everything, even converting products into services. Software and systems really do change radically because of the destructuring that takes place as you hook things onto the network. Ultimately this line between bits and atoms is going to blur. We'll get into this state, if you will, of infra-destructuring. And that is the end—I told you it would be different.

How did we get stuck in the back row?

Why can't we have the entire conference out here?

Last Mile Turmoil

monday

august 25

4:30 pm

Jay Adelson:

It's the end of the day. You've all been very patient. So this better be very exciting and controversial.

I am the founder and chief technical officer of Equinix. Raise your hand if you've ever heard of Equinix, and be honest. Okay, maybe a third of you. That's actually pretty good. Cutting edge crowd. Seriously, it was about three weeks ago that the Department of Homeland Security Oversight Committee and the Congress asked a number of companies to help them understand infrastructure, security awareness, and the whole cyber-security world. They asked several companies to testify: AT&T, Dell, Microsoft, and Sun. Equinix was also asked to testify. Why is that? It's because we are at this point the center of where most Internet traffic converges. Currently, for those of you who don't know, there are over 120 networks that interconnect at Equinix locations around the world, and it gives us visibility into what's going on in the Internet space, both on an economic level and a technical level. The last mile, you could argue, is not in our world because we're at the core of the Internet, and we obviously see the traffic as it passes through. Whenever you send an e-mail from an AOL to an Earthlink account, AOL and Earthlink will touch inside Equinix's locations—fifteen centers around the world—but those last mile providers, the ones that are delivering it to the eyeball are very much participants in this world. I'll explain that in a minute.

At this point, we're seeing that the players have changed and the way that money flows between them has changed significantly enough to impact the last mile. If you look at the traditional participants inside one of these Internet business exchanges, it used to be that in the economic hierarchy you had the guy in the middle that would charge, say, Yahoo to get to the DSL user, and the

*** CTO & Founder, Equinix**

DSL user would pay, and you had the guy in the middle charge you on both sides. That has changed significantly, and it's altered the economics of how much it costs the last mile provider to get you connected to the Internet. Now over 60 percent of the traffic from the big content providers passes directly to the end users. That's a large portion of the costs that just a year ago they had to incur by paying the Sprints, AT&Ts, UUNETs, and so on. This is a big change. Also internationally it's a big change.

Why is that important? If you consider the old problem with Internet growth and what the issues were keeping the Internet from growing to its ultimate potential, it used to be a core issue. How many of you have heard of MayEast? More people knew about Equinix—that's exciting! MayEast was one of the first access points where all the Internet companies converged. It is operated to this day by MCI. As the Internet grew, a core problem developed that kept people from getting quality of service across the Internet. It was congestion on the single switch in the middle of the magic Internet cloud somewhere. We had to address the congestion in the center. That's one of those fundamental things that's changed.

George talked about it at the Gilder/Forbes Storewidth Conference and in other discussions. If you look at the trends you see a number of participants taking matters into their own hands, including the Verizons, BellSouths, and SBCs. Instead of sitting behind other networks in a hierarchy, they're actually moving into these exchange points and exchanging traffic directly. The number of independent entities that now trade traffic has increased by 200 percent in the last twelve months. On the actual central switching fabrics, you don't see that big of an increase in traffic, but what's significant is every time a packet passes directly between two players. It's no longer being paid for, it's free. It was SBC that recently merged 65 autonomous networks through Ameritech, Pacific Bell, and all of their assets into a single network and backbone. Arguably, a very powerful network now but roughly 60 percent of their traffic, which has now moved, is free to them. At an IP level, it doesn't cost them any money to deliver that traffic, a very important change.

How many of you have heard of MayEast? More people knew about Equinix—that's exciting.

Each one of the participants inside one of these exchange points has an average number of twelve connections to other participants. Remember, it was just a few years ago when the average number was one. This isn't just the eyeball network; this is also content. In any Internet transaction, you have your DSL customer, your core infrastructure, and at the end you have content. Content used to be economically constrained because you had to pay to get your packets onto the Internet and technically constrained because it had to sit behind a local loop. Now content has moved into the core. Companies like Yahoo, Google, Microsoft, Amazon.com, or Equinix connect directly to Verizon, SBC, Sprint, or whomever, so that the cost to deliver their product to the end user is virtually gone. The scalability of what they need to deliver is virtually infinite, so that has changed a lot. That leaves the final bit pretty much on the last mile.

If you talk to a typical last mile provider, they will tell you a very interesting statistic. If you talk to Cablevision or Adelphia or Cox or Comcast, they will

tell you that 40 to 60 percent of their public Internet traffic is not, in fact, destined to some specific content site or to some tier-one network, as it's traditionally called. It's destined to another eyeball network. This is extremely important because if I can connect those networks directly to each other, what does that mean in terms of performance, in terms of quality, but most importantly what does it mean in terms of costs? Congestion at the core is gone. There are currently no trends or laws which indicate that there will be scalability issues in the next five to ten years at the core. It's dark fiber; it's cheap; I don't have to trench it; it's from one user to another in a cage; it's pretty simple.

If we've eliminated the content and enterprise limitations and core congestion issues, obviously the bottleneck is at the last mile. But it's not a technology issue. As you know, and we've heard here today, there are some really great technologies that can up last mile bandwidth significantly. I've seen trials of pushing 10 megabits over 8,000 feet of conditioned pair of a DSL loop. No problem. The technology is there, so what's the headache? The headache is the big players, the incumbents that haven't had the money and haven't had the time and haven't had the infrastructure. They have not deployed fast enough. It's a proliferation issue. It's how fast can these last mile providers deliver. I think we heard at this conference that the answer is not fast enough. The good news is that because all that cost I previously mentioned has been taken out of their operating world, they no longer have to pay some tier-one network provider to get to the rest of the world. Their costs of delivery of public Internet are lower. Some of that money is shifting into deployment. I wish I could tell you that it was a large enough chunk to speed it up so you'll see it soon, but we are already seeing that have an impact on the cable companies.

The other piece to this, of course, is once you move everything into the core and you have this level playing field inside of a marketplace, like an Internet business exchange, you get competition with very small players. I had a lot of conversations about wireless Internet Service Providers today with a number of you. It's irrelevant whether these wireless ISPs use Wi-Fi or CDMA. The point is that the small player is having a lot of success in a real market, even in a suburb of Washington, DC, whereas the traditional carriers have had a lot of trouble. That's because the small players move lightning fast and their scale is much smaller. This competition, I think, will drive those larger brand names into delivering alternative solutions quicker.

I've seen trials of pushing 10 megabits over 8,000 feet of conditioned pair of a DSL loop. No problem. The technology is there, so what's the headache?

Andrew Odlyzko:*

You may have noticed that the title on the agenda is "Last Mile Turmoil." The title of this slide is "First Mile Turmoil." Seems like a very small change, but it's indicative of a very fundamental, implicit assumption that the industry is making that I think is quite misleading. Namely, that the consumer sits at the end of a transmission link and you pump content to him. You pump movies; you pump music. Just think, for example, about Paul Jacobs's presentation this morning about what's happening in wireless. He was talking about people downloading ring tones and other kinds of things.

You can think of consumers as passive entities. However, it is very misleading. It is not where the value is. Think about the recent U.S. Northeast [August 2003] blackout. Did people complain because they couldn't download ring tones or because they couldn't watch their TV programs or download music? No, their main complaint was that their cell phones didn't work. They couldn't call their loved ones and were delighted when their wireline phones worked. Content is not king. Connectivity is much more important. That's what really matters to people.

It helps if you think of the first mile as the place where the consumer is in the center of the action. The content that people will be producing is likely to dominate. There is general growth in peer-to-peer. Too often in the public eye, peer-to-peer is synonymous with pirated music or video. Well, there certainly is a lot of that, but peer-to-peer is the future. Think back to Greg's presentation about all those devices being connected. What that means is that most of our traffic is going to be machine-to-machine. It will be much more balanced and symmetrical than just downloading. It will encourage us to think of the first mile, not of the last mile.

Another very important, deep, and implicit assumption in regard to first mile turmoil, which I'll discuss in more detail later, is something that we're going to hear about in many of the presentations regarding wireless. Think again about Paul Jacobs's presentation. He told us Qualcomm is thinking, along with the rest of the wireless industry, that the future is in data. Well, yes, it is in some stage, but what about all the opportunities in voice? As a matter of fact, what you learned today is that only a fraction of total voice traffic is going over wireless links, typically under a quarter for most countries, for which I have data. There are plenty of opportunities to exploit voice. It's one of those really amazing things that somehow the industry hasn't been paying attention to. The industry has been talking about mobile data access; it wasted nearly $100 billion dollars on those 3G licenses in Europe; it hasn't been thinking about voice.

Voice offers very interesting opportunities for implementing ways of providing differentiated services. How many people are happy with the quality of their cell phone? Anyone? OK, a few people

Too often in the public eye, peer-to-peer is regarded as synonymous with pirated music or video. Well, there certainly is a lot of that, but peer-to-peer is really the future.

* Director, Digital Technology Center, University of Minnesota & Internet Traffic Expert

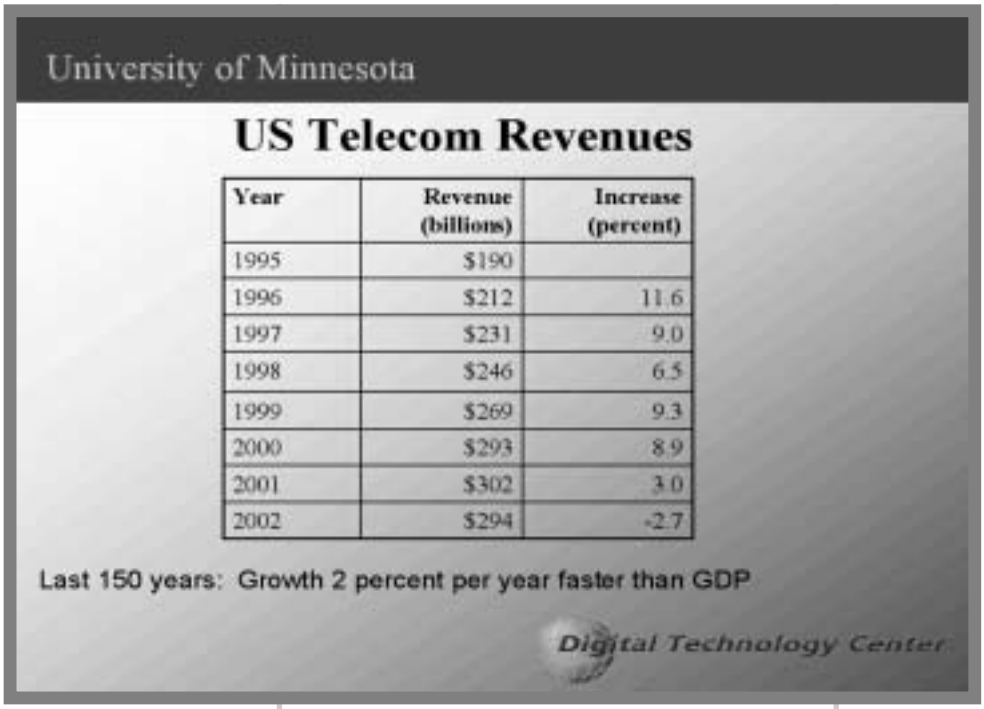

Alcatels, all have crashed, indeed. The service company sector as a whole has not crashed. Yes, growth has slowed down; as a matter of fact there's been a slight decline over the last couple of years, but that followed a period of very, very rapid growth. It was faster than regular growth and so it might just have been compensating for faster than expected growth. Historically over the last 150 years, U.S. telecom revenues grew about 2 percent faster than the economy as a whole. All the precedents show that's likely to continue over the long run, and there's probably a slight hiccup happening right now. But, of course, general stability consumes quite a lot. Many companies have crashed. We had bankruptcies of WorldCom, many of the IXCs (Internet exchange carriers), and many CLECs as well. Wireless carriers have grown; traditional IXCs have shrunk. Telecommunications is still very desirable. People pay a lot for it and are likely to pay more in the future.

are satisfied, but most people are not because it really is measurably inferior by design to that of wireline telephony. Voice offers opportunities for providing differentiated services with different tiers of services, acquiring more revenue, especially from business customers. There are all sorts of opportunities associated with voice that the industry is not taking advantage of. This is important for the first mile because what we are likely to see, and again there are some signs that the industry is finally tumbling into this opportunity and beginning to emphasize greater voice usage, is that eventually all voice telephony will be moving to wireless links. That's likely to promote a greater push into broadband by the wireline carriers.

There are all sorts of opportunities associated with voice that the industry is not taking advantage of.

But there's also turmoil throughout the industry. The most interesting things are in the first mile: the core is getting hollowed out.

The core of the network is not where the action is. The wastefulness of our investment in long-haul fiber networks accentuates something that occurred as a result of the developments in technology. Here's an example: 360networks built a transatlantic cable for $850 million. It was bought recently out of bankruptcy for $80 million, 2 cents on the dollar, in effect. The operating cost of that network is $10 million.

If you think of the telecom industry as a whole, it hasn't crashed. The supplier sector, the Ciscos, Nortels, Lucents,

The lit capacity was already in place at 192 gigabits per second. Eventually, it can be pushed into several terabits as you buy better equipment. Average transatlantic Internet traffic today is around 60 gigabits per second. In other words, the entire transatlantic Internet traffic could go over this one network that costs $10 million per year to operate. Suppose you pay this guy $50 million per year. That would make him a nice profit in any case. How much would it cost? Well, for the average residential user in the United States who typically downloads about 2 gigabytes per month per subscriber that would be a fraction of a dollar per month. And most people don't send much traffic across the transatlantic anyway. Long distance networks are not that expensive. Costs are very low for long-haul networks. Essentially all the action is either in the first or last mile, and it's very interesting because that's where the costs are. That's also where the revenue opportunities are. How that's going to be played out is not entirely clear, but wireless is becoming an increasingly disruptive influence.

Moore's law is improving all electronic technologies. It's certainly improving DSL as well as cable modem and wireless technologies, especially fixed-wireless technologies. But there's one big, big difference. For wireline technologies where they're talking about copper or co-ax or fiber to the home, there is a basic irreducible component of about $1,500 per household. If you want to take fiber to the home, not even counting the electronics, which are expensive but shrinking rapidly, it will cost you about $1,500 per household. That is primarily for regular labor, pulling the fiber, splicing the cables, etc., not for the electronics. No one can reduce it much further. You don't have the same issue with fixed wireless. Yes, for your access points you still have the cost of some number of thousands of dollars to install it, but if you can share a single access point with a hundred households in the area, then the cost per household becomes negligible. That's why there's so much excitement about wireless.

There are always interesting questions about whether 802.11 is the future. I'm sure it's not. But 802.11 is a forerunner of what we're likely to see. Fixed wireless, I believe, is likely to play an increasingly important role in pushing broadband to the home. It doesn't necessarily mean it has to dominate and provide most of the connectivity. The wireline carriers have connections that are already there. They can increase their capacity relatively quickly at low cost because the level of their electronics is improving. But they have to write-down the good value of their plant, which could result in lots of financial turmoil. However, they might hang in there; they might keep most of the excess market, but it's fixed wireless that will help push them in either direction. Plus the wireless substitution of voice: once the wireless industry really wakes up to the opportunities and moves 3G toward providing voice rather than data, then you'll see a lot of first-time migration of voice to wireless, creating more of a push to generate some broadband use of excess. I think we're likely to see much more turmoil over the next few years. It will be quite an exciting time because I foresee that we'll have a very rapid deployment of broadband connectivity and lots of opportunities ahead. Not everyone is likely to win, unfortunately. I have more details here in my papers if you're interested.

Essentially all the action is either in the first mile or the last.

"Do you think we should be taking notes at lunch too?"

Too cool for Telecosm.

Spencer Reiss and Rich Karlgaard critique the swings of golfers at the golf school down below.

"I hate getting stuck at the kids' table."

Metcalfe, Dayton et al. plot Wi-Fi's takeover.

Beyond the Technology

beyond the technology: observations on the economy

tuesday
august 26
8:00 am

George Gilder:

It's always a great joy to welcome Steve Forbes to this conference. He brings to it something that is very rare in any gathering of people, and that is true wisdom. He somehow absorbs all the information from all his sources and all his contacts and somehow distills it into a complete and balanced and fair and sometimes brutally frank exposition of his wisdom. That is what we uniquely have to hear this morning. Welcome, Steve Forbes.

Steve Forbes:*

On a downcast morning like this it's always nice to get flattering words. My only regret is that I didn't bring my fifteen-year-old daughter along. George talked about my having wisdom, but that sometimes isn't so clear to a fifteen-year-old.

It is a great pleasure to be here this morning, and I want to thank George and his colleagues for making this conference possible. I hope I don't say anything this morning that gives you indigestion, but I think this meeting turns out to be very, very timely because the economy itself is poised for a very significant recovery. Even though you'd never know it from the daily news—today's headline in *USA Today* indicates that casualties in Iraq since May have equaled those of the war itself and it looks as if things aren't going very well overseas, but the fact of the matter is we are—even though it's a cliché to say it—at the true proverbial crossroads. The question before us today is: are we going to go the way we did after the First World War or are we going to go the way we did after the Second World War? After the First World War—I want to give you some historic perspective because I think it's very, very important right now in terms of what our policymakers do and the kind of environment you're going to be operating in—huge mistakes

* President & CEO, Forbes Inc. • Editor in Chief, *Forbes* magazine

were made. The Versailles Treaty was a disaster; reparations poisoned the politics of Europe. In the late 1920s and early 1930s we passed the Smoot-Hawley Tariff, which was probably one of the most destructive pieces of legislation in history. It destroyed the global economic system, gave us the global depression, and was followed by the disastrous diplomacy of the 1930s.

After World War II, by contrast, we had one of the most creative, constructive eras in human history. We had the Marshall Plan, even though all that pouring of money didn't revive Europe. The institutions that were brought in at that time helped bring down the pre-war barriers, helped liberalize Europe economically, and showed that the United States was not turning its back on the world. It had its unique obligations. We established GAP, now the WTO, and began the whole train of free trade agreements, which was enormously influential in postwar prosperity. We had the Bretton Woods International Monetary System, and even though that formally collapsed in the '70s, the system itself in its aftermath prevented a re-descent to the purgatory of the beggar-thy-neighbor policies of the '20s and '30s. At the same time, we created NATO, which provided Europe with a sense of security. Even though the Russians were on the border, the fact that NATO was there gave people a sense of security and a sense that they could invest for the future. Even though huge mistakes were made in the Korean War, we showed that we were willing to shed blood to wage the Cold War. We talk about World War II and the "Greatest Generation" because we won this extraordinary conflict, but even more impressive we won the postwar peace, which we didn't do after the First World War. We're beginning to appreciate again, thanks to Iraq, how important it is, how difficult it is to win peace. Those policies did work. They de-Nazified Germany and allowed [Ludwig] Erhardt's reforms to go forward that brought about the German miracle. Conrad Adenauer got Germany off its mythical love affair with the East, which was very destructive for German politics, and firmly tied Germany to the West. Japan turned its back on ghastly militarism, and after an eight-year occupation democracy was firmly planted there. I emphasize that it did take a seven- to eight-year occupation to do it. All of our World War II enemies became liberal democracies. Europe had a swift recovery, and even though the early years after the war were pretty rough, by the early 1950s things were magnificently back on track again.

After World War II, by contrast, we had one of the most creative, constructive eras in human history.

Even though many policies responded to matters of the moment—we didn't know we were going to have a Marshall Plan; we didn't know how Bretton Woods would unfold; we didn't even have a conception of NATO in 1945—we did respond on the basis of these liberal democracy principles and largely it worked. Why this history? Because I mentioned that we're at the crossroads again. Although what we face today is nothing as destructive as what happened in the Second World War, and though the enemies we face today are not as powerful as the Soviet Union was in the 1940s and 1950s, the reality is we do have things we've got to get right, or we're going to have a much, much sadder and unsafe world. It's critical that we get things right econom-

ically and politically in this war against terror. If we do win it, and I'm going to touch on this in a moment—and this economic recovery is part and parcel of it—we will have the safest, most prosperous, and most liberal world order in human history. We have it within our grasp to do it.

In terms of energy, there are some sensible things we should do such as putting pressure on OPEC to lower oil prices. We have the whip hand; we should do it.

Today in terms of the economy itself, and I'll give you the usual caveats about the recovery—if the Federal Reserve goes off the deep end, if the Korean Peninsula blows up—if we don't have some 9/11-like disaster, this recovery is for real. Productivity is real. Companies as we all know so painfully are much leaner than they were three years ago so operating profits are improving. Even in IT there are some signs of new life. Durable orders are good; retail sales have been remarkably good, although consumers still feel lousy. The worries over the rise and jump in interest rates that worried economists a couple of months ago, I think, were overblown. Interest rates today are still about where they were a year ago, which in turn are lower than they've been in the last forty years. We haven't seen rates as low as they are today, so that's not going to abort the recovery. In terms of energy, there are

some sensible things we should do, such as putting pressure on OPEC to lower oil prices. We have the whip hand; we should do it. We haven't and I don't know why we haven't. I don't know why we haven't played with the petroleum reserve. There are lots of things that can be done. That's why we've got to get Iraq back on its feet. We can break the back of OPEC, and I hope in the next couple of years we do so we get true market prices in that area again.

You keep reading in the paper how much tax cuts cost. It's a junk number. Talk about junk science, this is the ultimate example of it.

There are two major factors that I believe will make this recovery real, given we don't have some international disaster. One is the tax cut bill that was passed a couple of months ago. The only thing you have to understand, bottom line, about a tax bill—put all the garbage you read in the papers every day aside—is whether it reduces the tax rate that you pay on each extra dollar of income. What the economists with the usual sense of poetry call marginal tax rates. If the marginal tax rate is the tax rate you pay on your last dollar of income, and on your next dollar of income that rate is lower, the tax cut is good. If the rate is the same, waste of time. If it's higher, bad. That's all you need to know about tax policy. The rest is garbage. The thing to keep in mind is taxes are not just a means of getting government revenue; they're a price and a burden. George knows this; many of you know it instinctively. Taxes are a price and a burden. The tax you pay in income is the price you pay for working. The tax you pay on capital gains, if you have any, or on profits, if you have any, is the price you pay for being productive, successful, and willing to take risks that work out. And the proposition is very simple. When you lower the burden and price on good things like productive work, success, and risk taking, guess what? You get more of them. Very simple, but the political culture finds it impossible to grasp it and understand it. Actually, for once, the sausage factory in Washington got this tax bill right, largely. Unlike the tax cut of 2001, which was phased in over 500 years and was useless, this one really did some things right. Most significantly, it reduced the capital gains tax of 20 to 15 percent. The reason that's underappreciated, as you well know, is that it helps those who are not in the daily headlines. It helps the startups that we don't even know about but will in ten or twenty years down the road, so that was a very, very positive thing. And that almost happened by happenstance. The House insisted on it because of a few members of the Ways and Means Committee; the White House was cool to it and had been cool to it in 2001. Somehow it got in. Miracles occasionally happen in Washington.

Reducing the dividend tax: Bush wanted to eliminate it; instead of getting the whole elimination, he only got two-thirds of it. But the lesson there is if you ask for the whole loaf or maybe two loaves, you might get two-thirds of a loaf. If you don't ask for it, you're not going to get it. If Bush had asked for, say, a 10 percent cut in dividend tax rates, he'd have gotten zip, zero. Instead he got 60-plus percent. I wish they'd keep that in mind when they negotiate, as they will in a few weeks, on the Medicare prescription bill. The White House has made it clear that they'll just sign anything that Ted Kennedy and others shove at them. If they'd take the

same kind of attitude that they did on the tax bill, they could actually give us something on prescription drugs that wouldn't be a long-term disaster, that would do far more good than harm. On the tax front, they did things right. You keep reading in the paper how much tax cuts cost. It's a junk number. Talk about junk science, this is the ultimate example of it. Every time you reduce tax rates, the marginal tax rate, the economy gets stronger and government revenues end up going up not down, without exception. The idea that this bill is going to cost us money over ten years—as George said they can't predict weather for tomorrow—we don't even know what the economy will be like in ten years. That ten-year estimate of what a tax cut would cost is a bogus concept. Tax cuts properly structured create a stronger economy, but the numbers are meaningless. It's about as meaningful as your adding up all the weight you've lost on diets in the last ten years. It can be a pretty scary number, but as some of us amply demonstrate it's not a very meaningful number. That's about what these tax cut estimates come down to.

Already, as George and others have pointed out, since the tax cut passed in May, the increase you've seen in stocks alone is why we have such a good bull market. People are very unhappy—no one is as ginned up as they were in the late 1990s—and that's a good bullish sign. Bull markets always climb walls of worry. That's a good thing. The fact you're sad is a good thing for the market. In terms of the market itself, it's already gone up over a trillion dollars since the tax cut passed in May. If you've ever looked at a balance sheet, and John Rutledge is one of those few people who do, this tax cut has already paid for itself.

As for the flat tax, one of these days even the United States is going to get it. As you know, Russia did it two and a half years ago, putting in a 13 percent rate, even lower than mine. I never thought Vladmir Putin, Communist-minded president of Russia today, former Secret Police officer, would ever become more radical than I on taxes. My rate was 17; he did it at 13. It's been a roaring success. Here's a country with no tradition of paying income taxes, and in the last two and a half years the flat tax has really gotten people accustomed to paying taxes because it's low, simple, and easy to understand. Now they're cutting corporate rates to lower levels as well. Even before the Russians did it, the Latvians and the Estonians did it—the

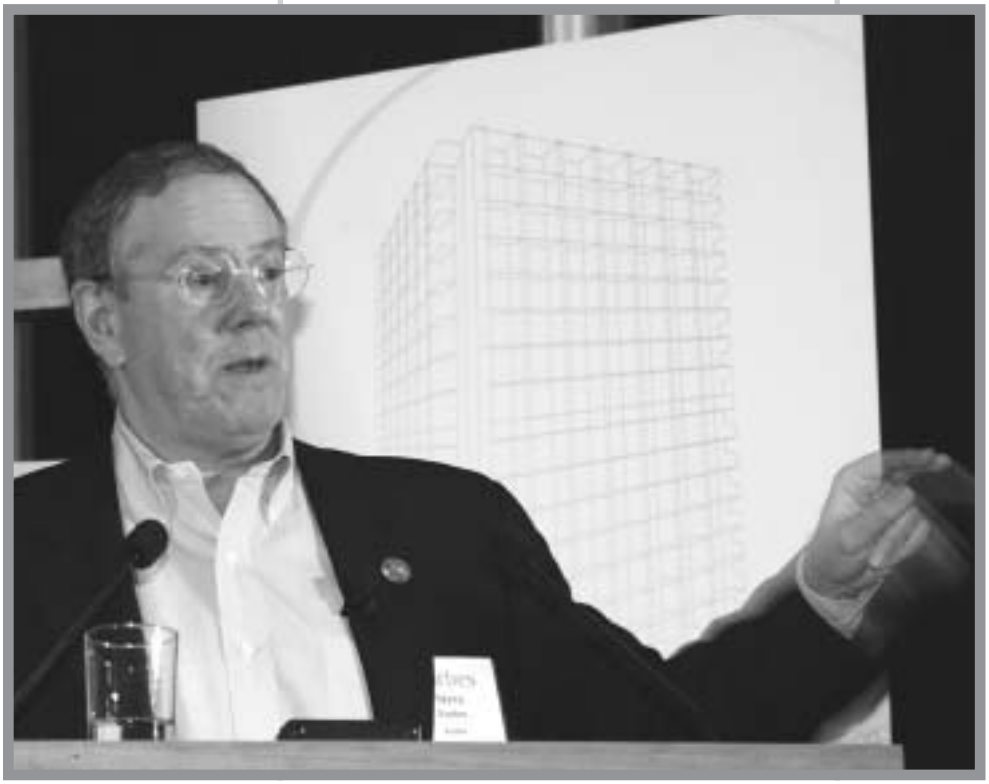

Baltic States. Because they're small, nobody pays attention to them, but it's worked. Ukraine just passed it recently. If they ever get rid of the corruption in that country, its economy should start to show signs of life. And Slovakia is about to pass it. The reason little Slovakia is important is that Slovakia is joining the European Union next May. They're part of the ten new members coming in. That's a good thing because the Slovaks actually put in its flat tax at a 19 percent rate. Guess what? The Czechs will have to do it; the Poles have indicated they'll start to do it. You'll begin to get a ripple effect. These new members of the European Union have no interest in being museums, of having giant welfare states that make their countries tourist museums for people like me to visit in order to see what civilization looked like a hundred years ago. The Eastern-Central Europeans want growth, and I think what Slovakia is doing will be a precursor of a change in attitude in the European Union. So on the tax front there's actually been some good news.

The other piece of good news from a very unlikely source is the Federal Reserve. The Federal Reserve has put us through a deflation since the late 1990s. Unfortunately, economists have about as much understanding about deflation as doctors did about diseases 200 years ago. They're clueless. We haven't had any kind of real deflation, sustained deflation, since the 1920s and 1930s, and because it's new they're not familiar with it. Just remember how long it took for us to understand inflation in the 1960s and 1970s and to get a grasp on that. The Fed is clueless. You have the economic equivalent of dehydration. The Federal Reserve was inadvertently starving the economy for liquidity. Here's the reason why they didn't realize it then and still don't realize it today. Think of it this way: what would happen if you were told there was going to be a water shortage? You know what would happen. You'd buy all the bottled water you could; you'd hoard the water; you'd be very careful about using it. Alan Greenspan would then visit your house, see all of this water there, and say, "What liquidity shortage?" What you saw happening here was that money funds were taking in hundreds of billions of dollars.

Now the Federal Reserve is getting it right, but not because of any wisdom. Alan Greenspan, the great genius that he is, is still flying by the seat of his ample pants.

Economists looked, didn't understand the disease, and wondered how there could be a liquidity shortage when people were sitting on all this money. You could see it unfold in Japan, which has had over a twelve-year deflation. Bringing industry rates down is not enough. You can't bring down the nominal cost of money; you have to make the money available. It's like going to a gas pump and being told gas is fifty cents a gallon and then told you can't buy any. As John Rutledge indicated, we did have a credit crunch in the last couple of years, exacerbated by bank regulators. Japan, where they had virtually zero percent interest rates, has had a twelve-year recession. They're finally beginning to show signs of life that they're climbing out of it. But when the Fed put us in this mild deflation in the late 1990s, it first hit commodity prices, which had political repercussions in that disastrous Farm Bill. It hit steel prices, which is why we got the pressure for protectionism. The political fallout from deflation devastated countries in Latin America, hurt Japan, and made deflation even worse. It had a ripple effect and we're absolutely oblivious to it. That's why you read in papers that companies lost pricing power, couldn't raise prices anymore. You were under relentless pressure with expenses, the opposite of inflation.

Greenspan, and it's true of other central bankers, looks as if he's always coming from a funeral or memorial service.

Inflation first gives you a high because the money comes in, and then your expenses catch up and then some. With deflation, it's the opposite. Your revenues get hit; you're under relentless pressure to cut costs; you don't quite know what's happening. This is why the so-called bubble in hi-tech was so bad, why it was so much worse than what we had in the automobile industry in the last century, or railroads in the nineteenth century, or personal computers in the early 1980s where we had a big shakeout. This time it was actually made worse by the deflation. It wasn't just a big investment boom followed by a big shakeout. That was part of it, but it was also made far worse by the deflationary environment, which first hyped investment and hi-tech and then slammed it when the deflation really began to bite. It was devastating. There is no research being done on it, and most people are absolutely oblivious as to what happened to them.

Now, finally, the Federal Reserve is getting it right but not because of any wisdom, and that's the danger. Alan Greenspan, the great genius that he is, is still flying by the seat of his ample pants. The Federal Reserve doesn't know each day whether it's getting it right or wrong. It's like flying without instruments or driving a car without a speedometer or a fuel gauge. Greenspan has some good instincts so he doesn't get off course too much, but the Fed still doesn't know what it's doing. That's why you have uncertainty in the markets.

I've given you advice on how to be tax experts. The only thing you have to look at in terms of monetary policy, you busy people, is to look at the price of gold. It's the best barometer out there. Commodities are the most sensitive barometers of monetary disturbances. If they're all going in one direction—all up, all down—watch out! Something's not right. But the easiest way is to look at the price of gold. Right now it's about $350 to $360 an ounce. That's good. If it goes above $400, bad. If it goes down to $300 or below, bad. That's all you need to know about monetary policy,

and you're better than every other central banker in the world is. You really are. As long as the price per ounce stays in the $350 to $360 range, that liquidity, that hoarding of dollars, will start to seep into the economy. Even though it's going to take longer because we still have the drought mentality and don't know it, the fact of the matter is the liquidity will be there if the Fed keeps it at that range.

This is also a moot issue. Economists love the idea of manipulating currencies. They think that shows sophistication: we can get better terms of trade, this, that, and blah, blah, blah. Actually, it's grossly immoral. If you work and receive a dollar for your labor, why should politicians and central bankers determine whether you should get extra for your labor or cheat you out of what you are awarded for your labor? It should be a standard measure of value. If your labor is worth a dollar, the politicians shouldn't say you only get ninety cents or a bonus of ten cents for a dollar ten; it should be a standard measure of value. It's like a watch: it's sixty minutes an hour, last I looked. Imagine how difficult our lives would be if we floated clocks and watches at sixty minutes an hour one day, forty-eight the next, ninety-six the next. Economists would say that if we increased the hour from sixty to sixty-three minutes, by golly we could get extra labor from our workforce and increase profits. It's very pernicious nonsense. Immoral—don't do it.

The other thing to watch out for with our friends in the Federal Reserve is that there's a gene for central bankers called the "sourpuss gene." Greenspan, and it's true of other central bankers, looks as if he's always coming from a funeral or memorial service. They always look as if they've eaten something that didn't quite digest right. The danger is that if this recovery does begin to gather steam, especially after the elections next year, the Fed may start to tighten because it thinks we're overheating again like the late 1990s. That's a danger. The sourpuss gene is always there, and when people like you are happy the Fed gets very, very antsy. I don't know why, but it does.

What's happened with the trial lawyers and the pernicious, barbaric ways they have distorted the legal system means we don't have a legal system that people can count on.

Two good things have happened on the tax front and the Federal Reserve front. Let me quickly give you some things to watch out for, things that can go wrong. I mentioned the Federal Reserve; if gold stays around $350 to $360 an ounce, rest easy. Taxes—just look at tax rates. One of the wise things Californians did a few years ago was pass a referendum saying that before the legislature can increase taxes, it needs a two-thirds vote in the legislature, not a simple majority. If it had been a simple majority, the budget crisis would have been made infinitely worse because they would have raised income tax rates, and that would have been a true disaster. California is bad enough, so thank God for that. Watch out for tax rates.

Another thing that is the cancer of capitalism is the disaster in our legal system. I'm not going to go into any detail about that this morning. But what's happened with the trial lawyers and the pernicious, barbaric ways they have distorted the legal system means that we don't have a legal system that is reliable, reasonable, that people can count on, that if you do something you're going to get a certain result. It is truly like a lottery. People who really suffered from asbestos, for

example, in World War II in the shipyards may have gotten $8,000, $10,000, $20,000, $30,000. Peter Angelos, owner of the Baltimore Orioles, got over $100 million. It's a disastrous system and one that has really warped the way that we look at the world. Everything we do now is always motivated by or certainly restrained by legal fear. You coach your soccer team or Little League team. You always ask yourself what's my liability? A teacher disciplines a student. What's his liability? You give a party—swimming pool—what's my liability? If you *say* something, what's your liability? This fear of lawsuits is really a cancerous thing; it's eating at our innards. Fortunately, counterattacks are beginning to occur against this juggernaut, this assault on civilization. The Chamber of Commerce is doing it; local businesses are beginning to do it. Texas, which was a haven for these vultures, passed a very good reform law a few months ago. Mississippi, an even bigger haven for this kind of corruption, passed a pretend reform bill. Even there some of the judges are beginning to behave themselves. Why? Because of the absolute corruption of the trial lawyers who began to bribe jurors, give loans to judges, things like that.

This fear of lawsuits is really a cancerous thing; it's eating at our innards. Fortunately, counterattacks are beginning to occur against this assault on civilization.

I'll give you one example. A friend of mine told me about a friend of his. He was a defense lawyer, a good lawyer, the right kind, and he was going to have to spend the summer in Mississippi with some frivolous lawsuit. His family was mad at him because they couldn't go on vacation; their Dad was going to be in Mississippi defending this garbage. Suddenly in June the lawsuit was dismissed by a judge who had always given out outlandish awards. Why? Because

down the block the Feds had a corruption investigation of his suborning juries, bribing judges. This guy wanted to show that he could render objective decisions, so the idea of jail concentrates the mind wonderfully. Change will come eventually, but it's going to be a relentless state-by-state fight. You've got to get involved with it and help that fight. Contact your Chamber of Commerce because this is a cancer that has to be fought.

Another thing to watch out for is the aftermath of Sarbanes-Oxley. There have been some good reforms on governance, but one of the dangers—and it's a real one—is confusing genuine fraud, genuine wrongdoing, with risks that don't work out. You can appreciate better than most that risk taking is the essence of free market, free enterprise, democratic capitalism. Most new businesses don't work. Most ventures don't succeed. Even established companies go into ventures that fail. The worst thing that can happen is if we have a system where if you don't succeed you're subject to criminal and civil penalties. One of the great engines of advancing Western civilization was the creation of the limited liability corporation of Britain. If you put a certain amount of money in a company, that's all you're liable for unless you've committed a true crime. You didn't go to debtor's prison if your business failed. If you lost your money, that was it, you could start again. If limited liability gets undermined, forget it. We're not going to have the high standard of living we can achieve. We're not going to develop a lot of the things that are coming down the road in terms of new advances and technology if people are always looking over their shoulders wondering how things will look in a courtroom. Every new idea looks like it was a given that it would succeed when it succeeds. But every new idea looks pretty dumb when it doesn't succeed. You don't know until you try it. Who could have explained the concept of the automobile 120 years ago? It would have been impossible. Who could have explained the Internet to the average person twenty years ago? Impossible. If you do a spell check from 1988, type in Internet, guess what? No such word. The advances we get come from people doing seemingly outlandish, dumb things. When they work, it looks like it was preordained. When they don't work, you ask how could they have been so dumb. They must have had criminal intent. They hyped the thing. They said this was better than the wheel. This new bread was better than the old bread. You always hype when you think you have something new because you're trying to sell it. You're not going to go to a venture capitalist and say, "I don't know if this is going to work because I've got this and that . . . maybe it will, but I don't know." Unreal. If you allow it, technology does move on, even in the aftermath of what happened in the late 1990s and early 2000. But the thing to watch out for in Sarbanes-Oxley is that we don't allow the courts and regulators to start making it a crime to fail. If you do that, it's over.

The thing to watch out for in Sarbanes-Oxley is that we don't allow the courts and regulators to start making it a crime to fail. If you do, it's over.

Another thing to watch out for, as George and others have pointed out, is regulation—what we should do with the Federal Communications Commission, for example. I don't have to tell you—the Internet has obliterated all the traditional differences we have had between telephones, TVs, long distance

and local telephony, etc. Technology has obliterated all that, and yet the regulators still pretend we live in a world where we can precisely compartmentalize everything. It's nonsense. What should be done in communications is that they should pick a date, like they did when they got rid of fixed commissions in the mid-1970s in Wall Street. They said by date certain all the price fixing is over; you're on your own. Pick a date, twelve or eighteen months from now, and deem that all the old barriers go away. If you want to be in satellites, if you want to be in local, long distance, whatever, you can. No barriers, just do it. Give people time to adjust to the new world and let the market, i.e., people, determine what emerges from it. What we should do at the FCC is send it to North Korea and let them undermine what's left of that country. Watch what happens in the regulatory front there. As George has pointed out to you, a typical household there has the equivalent of four T1 lines, which would cost you $3,000 a month in this country, for $39.00 a month. George tells us that Italy is starting to make moves in terms of really opening up broadband. We're in the Dark Ages on that, and the sooner we get it opened up the happier conferences like this will be. I mentioned Medicare. If we create a great new entitlement there, it's just going to be a bigger mess down the road. Watch what happens on that front.

What we should do at the FCC is send it to North Korea and let them undermine what's left of that country. Watch what happens in the regulatory front there.

Right now things are beginning to move in the right direction. Even overseas, in the rest of the world that practices economics in a peculiar way, things are starting to go in the right direction. You look at Japan. Our reflating, which we've been doing now for about eight months, is starting to bring life back to that country again. And who knows if this new head of Bank of Japan will really reflate that economy, but the stock market has gone up about 20 to 25 percent in recent weeks. It looks like there are signs of life in Japan. China has huge political problems, but they still seem to be moving forward. Europe, Old Europe, even there we find signs of life. They lag us; it won't be until next year that they'll get the kind of recovery we're starting to experience now. Even Germany is in the throes of reducing tax rates—not much—from 48 down to 42 percent, but at least they know they have to begin

to make reforms. Even the French know they have to make reforms in pensions and things like that. So it's going to be slow, but at least it's moving in the right direction instead of profoundly in the wrong direction. That's why things are beginning to look up.

In terms of the International Monetary Fund, which is the most destructive agency out there, they practice economics the way doctors practiced medicine 200 years ago. When you got sick, they would bleed you. The IMF does the same thing today. They bleed you to death with high taxes and devaluations. The IMF, for the moment, is destroying Brazil and what's left of Argentina. They're trying to wreck Turkey, but not as badly as they did a couple of years ago. The IMF is a destructive agency and the sooner we reform it the better. This has political consequences. I mentioned Turkey. Turkey, twelve years ago, was our firm ally in the first Gulf War. This time they profoundly were not. The reason is they had a change of government. They have an Islamic government in Turkey. How did that happen? Well, two years ago Turkey had a currency collapse, with help engineered by the IMF. When I visited Turkey as a teenager in the mid-1960s, one dollar bought you 60 Turkish lira. Today a dollar gets you about 1.5 to 2 million lira. Currency collapsed—guess what?—desperate people turn out the ins and put in people they never would have thought of because things are bad and they want something better. The IMF needs reforms. I think I've told you before what the reforms should be, other than getting rid of it. But, agencies are immortal; you never get rid of them. If you work for the IMF you're paid in dollars. Even if you're not a foreign national, you get paid in U.S. dollars and the IMF pays your income taxes. If you owe income taxes, the IMF writes the check to the IRS for you. In the future a great reform would be to have the IMF people paid not in dollars, but in the currencies of the countries they advise, (applause), and pay the tax rates of the countries they advise. It would change them overnight.

In closing, I'd like to mention something about Iraq. This is truly where the center between civilization and the force of darkness is having its match. Whether we should have gone to war in Iraq or not, or Afghanistan or not—I talked to a friend of mine the other day who used to work for the State Department and he said those debates are over—we are there; they are ours. He said if we don't do it right, there will be severe and savage repercussions, especially with Iraq. We're there; it's ours; if we don't get it right, that neighborhood is going to be a disaster. The terrorists will have new resources because they will overthrow the authoritarian regimes in that part of the world if we don't succeed. We are committed.

We made two big mistakes after the big war earlier this year. One was we didn't move quickly to decapitate the Baath Party; we're now starting to do it, but we didn't do it very fast, other than that deck of cards. The second thing we didn't do was rely on local Iraqis to do things they should have been doing like policing and guarding. Now we're beginning to do it. We're starting to train police and militia and that should have been done months ago. Why we didn't do it, I don't know. We can fight

If we are steadfast we can create an environment where our values can take root in parts of the world where they've never taken root before.

these terrorists if we work at it. We all know what happened to the UN Headquarters there. Our people had gone to the UN Headquarters weeks ago and said, "You are very vulnerable." One of the guys who works for the UN actually used to work for the U.S. Army and, almost surreptitiously, put up a wall. They couldn't put up the kind of wall that should have been put up as a barrier, but it was enough to prevent what would have been a carnage involving not twenty dead, but hundreds of dead. But the UN did not put up the proper barriers: they said you're Americans, we're the UN; the terrorists know we're the good guys; you're the bad guys. But we can fight this terrorism if we pull in the resources and stop doing dumb things. Do you know that Iraq is still using currency with Saddam's picture? Can you imagine using currency in postwar Germany with Hitler's picture on it? You just wonder what goes through these people's minds. However, they're finally starting to lurch in the right direction. The danger is that we will tire of this effort. We won't spend the necessary years there to get this thing done right. But if we do spend the time and get it right, we need to follow free enterprise policies, which we're not doing in Iraq now. They're not putting in the kind of liberal economic policies there they should. They're making a hash of the currency issue. I'll give you one example. This October they're finally going to allow people to exchange their Saddam currency for a new Iraqi currency, which is fine. But they should put restrictions on how much of the old currency you can turn in for the new, because who has most of

the old currency? Saddam, his alllies, his family, and his henchmen. In Nicaragua, we allowed the Sandinistas to turn in all their ill-gotten gains for new currency, in effect laundered the money for them. Should we do the same for Saddam's family's money? Launder that for them this October? They're clueless about it. They haven't even thought about it.

There are a lot of things that can go right in Iraq if we just follow sensible policies, and eventually I think we will. The bottom line, although we have some very real and dangerous situations with Iran and Korea on nukes and though Iraq has gotten off to a shaky start, if we are steadfast we can create an environment where our values can take root in parts of the world where they've never taken root before. While that might sound Pollyannish and outlandish, it's no more outlandish than a few fifty years ago when it was said that Europe today would be a continent where war between the major powers was absolutely inconceivable. For a thousand years, Europe had been a cockpit of one war after another. During the twentieth century the two most destructive wars in history came out of Europe. The most destructive ideologies in the world, Nazism and Communism, came out of Europe.

Look at the Middle East today. There's a book called *Terror and Liberalism* by Paul Berman that outlines how this so-called Islamic ideology comes right out of Nazism and Communism. The Baath Party ideology is a pastiche of hideous European ideologies mixed in with some local stuff. Europe has been a source not only of great civilization but also of some really disastrous ideas and events. Who fifty years ago would have thought today that Europe would think it a big deal to veto a UN resolution? That to them was a big thing against America: we're going to veto the resolution and give them the middle finger. Throwing spitballs at us. They meet every few months and give papers to each other and try to put this union together. Europe today is very different from what it was fifty or sixty years ago. Even though it's a pain in the neck, Europe today is a success of American diplomacy. What had been a cockpit of power politics, bloodshed, bad ideologies today is nonentities. As I say, they think it's a big deal to veto a resolution.

The American economy is starting to show real signs of life, if we just but let it do so.

That's very different from what the Islamic fanatics want to do to us. Imagine the Middle East fifty years from now where Iraq, Egypt, Algeria, and the new Palestinian state meet every few months, pass resolutions, and think it's a big deal to veto a UN resolution against the United States. We'd say, "That's pretty good compared to what it is today." It can happen.

In sum, the American economy is starting to show real signs of life, if we just but let it do so. Globally, make no mistake, these next few years are going to be as critical to the future of this world as the late 1940s and early 1950s were after World War II and the early 1920s and late 1920s were after World War I. The only question is: are we going to go the World War I way or the World War II way? I think we will lurch fitfully in the right direction to more World War II postwar policies than the disasters of World War I. If we do, we're going to have many, many happy conferences in the future.

I'm quite happy that Davis and his ilk are getting their comeuppance. They really did enormous harm to California, and they should pay the price for it.

Gil Amelio:*

Could you comment on the situation we have here in California?

Steve Forbes:

Well, the real circus in California came before the recall. Here you have the biggest state in the Union—a state that's been a magnet to Americans for 150 years; the state that has great universities, great intellectual resources—and the political class has done its best to destroy it. You have a legislature where most of the seats, like Congress, are non-competitive. You have the Democratic Party, sometimes connived with Republicans, that is truly tax and spend. They repeatedly take the good times for granted and keep doing things that are going to harm this state in the future. In the late 1980s, we did a cover story in *Forbes* saying the good times are coming to an end. It was not well received in the Golden State, but it was true. When the recession hit, California was hit hardest. Remember Pete Wilson who implemented big tax increases and the big debates in the state legislature when they increased the sales tax? They exempted food, but were crackers and peanuts food or snacks? Should they have a snack tax? There was a crazy budget crisis and California was slow to recover. For a while it lagged the nation. We forgot that in the boom times of the late 1990s. Sacramento forgot it. California is very dependent, peculiarly dependent on options and capital gains. It should have learned from the 1980s and did not and pretended the market was going to go up 30 percent per year, forever. This latest budget crisis was even worse than the one in the late 1980s and early 1990s. So the people of California had a mechanism where they could say that this nonsense has got to stop. Sociologists tell us that every organization becomes self-centered, more interested in itself than its mission. You see that in the private sector; you saw it in the auto industry in the 1960s and 1970s until they got a wake-up call from the Japanese and other competitors. You see it in nonprofits all the time.

* Senior Partner, Sienna Ventures

They become self-contained and oblivious to what their true mission is. It happens in the political culture as well. In business, if you get self-contained too long, your customers, the marketplace, will upend you. In politics in California, it's the recall that is upending politics as usual. Although Republicans initially were opposed to the recall because they wanted Davis to stew in his own juice for three years and then take over in 2006 and get the benefit in 2004, the people of California said that they weren't interested in these little Beltway calculations. We have a disaster on our hands and we're going to do something about it. I'm quite happy that Davis and his ilk are getting their comeuppance. They really did enormous harm to California, and they should pay the price for it.

Will Arnold do better? Well, I took heart from his press conference the other day when Warren "Buf-fay" or Warren Buffett—a great stock picker, brilliant man, who I don't want to confuse with Jimmy Buffett—piped up that he thinks California needs higher property taxes, and Arnold was pretty good about it. I won't imitate his accent, but he said that if Warren opens his mouth again on taxes, he's going to have to do 500 pushups. And then Arnold twisted the knife and said, "And Warren, you're in no shape to do that." So a guy who can react that quickly when he has a political storm on his hands gives me hope that anything he does will be better than what Davis and his like did. I'm glad that this circus is happening because the real disastrous circus happened before. California's budget crisis, as severe as it is—here's a state with a GDP of about a trillion and a half to two trillion dollars and assets about four times that—a $30 billion deficit should be fairly easy to overcome if you put in reforms like reducing the capital gains tax, reforming an idiotic workers' comp system, rescinding crazy things like six weeks' paid leave, etc. Small businesses cannot afford that, much less even larger ones. Get rid of that kind of stuff and put in a series of sensible reforms and California will come back very quickly.

Since the late 1990s, we were misled by the boom in California to the real underlying crisis. For over 150 years, California has been a magnet for people. In the '90s, that changed. California still attracted immigrants from Mexico and Latin America, but other than immigrants coming from those areas California was losing population. California lost 600,000 people. That's the first time in 150 years that there's been an outward migration of Americans leaving California instead of coming into California. That should have told you that beneath what was happening in Silicon Valley this state was starting to have some real problems. Politically the reforms are going to be difficult. But if you do them, by golly, this state will come back very quickly. It wouldn't take much to turn this state around if you reformed the tax system, workers' comp, and some of the other dumb things they've done in recent years. People are drawn to California.

It wouldn't take much to turn this state around if you reformed the tax system, workers' comp, and some of the other dumb things they've done in recent years.

Mort Feldman (ATTENDEE):

If the G7 nations' currencies continue to strengthen against the dollar and let's say the euro goes to 130 from the current

108, wouldn't gold strengthen proportionately to where it would appreciate perhaps 15 or 20 percent with no effect on our economy?

Steve Forbes:

The question is if currencies like the euro or the yen appreciate against the dollar, won't gold do the same thing and would that be a bad thing? The answer is if we're doing our monetary policy right, gold will still be $350 to $360 an ounce. In euros it might be different. In yen it might be different. If the euro went from 108 to 130, Europe would be on the path that Japan has been on since 1989, a very severe deflationary path. In short, the supposed weakening of the dollar this year was not so much a weakening of the dollar as the dollar finally getting over the deflation of the late '90s and early part of this new century. There was a scarcity; there was a dehydration. That is now ending. As a natural thing there are more plentiful dollars, which is a good thing, so that other currencies would get different values against it. If Europe and Japan had their monetary policies right, you would not see much fluctuation in the price of gold, vis-à-vis their currency. The key thing is we can do monetary policy that keeps the dollar price of gold fairly steady. What other countries do, and they may go all over the map, is their problem not ours.

The supposed weakening of the dollar this year was not so much a weakening of the dollar as the dollar finally getting over the deflation of the late '90s and early part of this new century.

I see we've come upon the hour. Let me say, again, thank you for coming to our conference. Thank you for your very real kindness to me this morning.

Alex Lightman of
Charmed Technology

Andrew Raguskus of
Sonic Innovations

Latif Ladid of
the IPv6 Forum

A Sea Change in Semiconductors

silicon advances propelling the telecosm

tuesday

august 26

8:50 am

George Gilder:

Thank you, Steve. That was terrific. Do you see what I mean? That's how I get my guidance to the world.

I am intrigued by laws of various kinds and, I dare say, I occasionally abuse the privilege of promulgating them. I always liked Drucker's law; try to create jobs you get stagnation; go for profits and you get both jobs and prosperity. We all know and celebrate Metcalfe's law, which is the heart of the Telecosm, the exponential benefits of increasing connectivity. I also like Tredennick's law: seek performance in semiconductors and you do not get volume; seek volume and you get performance. Nick learned this in the semiconductor industry, designing microprocessors. The secret of the tremendous success of the microprocessor was the volumes it could attain because each one could be programmed to do a wide variety of functions. That's how the microprocessor became the dominant product. But Nick has always felt that the microprocessor has serious flaws in its separation of program from function. He has been the leading advocate in the industry of programmable logic where you get the performance benefits of hardware and the programmable benefits of the microprocessor and combine them to get both volume and performance. This is the paradigm that he has long expounded. He was chief scientist at Altera, which was one of the great companies pioneering this technology, and now he's going to explain it to us as he does regularly in the *Gilder Technology Report.*

Nick Tredennick:*

George has concluded my talk for me so now all I have to do is get through these stupid slides and we'll be ready for the break. When I put this presentation together, I had nearly 85 slides and since

* Editor, *Gilder Technology Report*

there was no way I could possibly get through them in the amount of time I had, I threw out about half of them. Unfortunately, I don't know which half I threw out.

We're in for a sea change in semiconductors that's as important as the introduction of the microprocessor was because today's components are unsuitable for where we're going. The microprocessor isn't good enough. In my opinion, digital signal processors (DSPs) are circling the drain, even though right now that's the fastest growing segment of the microprocessor business and Texas Instruments is here somewhere to defend it. Nevertheless, I wrote a paper a few years ago titled, "Death to DSP," and I think we're still going to see that. Programmable logic devices are unsuitable for where we're headed. So are all the memory and storage devices. It's an inevitable transition, but it's going to be extremely hard to do.

The microprocessor isn't good enough. Digital signal processors are circling the drain.

Here's the journey. In the beginning we did this thing with a PC. Now we've come to the value PC and the value transistor. The fabs are good enough today. Once the PC becomes good enough, the whole business is going to transition from what the PC was, a cost-performance system, to something that's an unwired, untethered system that changes the rules to cost performance per watt. It's a huge transition that microprocessors and DSPs just can't make.

The transistor was the first generic semiconductor. You could build anything out of transistors. You take that for granted. The integrated circuit kicked the industry into high gear. It kicked Moore's law into high gear. Moore's law isn't really a law; it's just how fast the industry runs the treadmill. You can either get more transistors on a single chip with each new process generation, or each transistor that you already have gets a lot cheaper. You build transistors, in two process generations, on these things like cookies that are big wafers of transistors. If you cut the width of the transistors, you're making transistors smaller all the time. Two generations will do this for you economically. If you go from 37 die per wafer to 180 die, you get six times as many chips. Over time, we shrink the size of those transistors. Something like five processes ago, we went through the wavelength of visible light. That means that we're building chips with transistors that are so small you can't even see them with an optical microscope. Not only that, but we've gone through so many generations now that we're building 90-nanometer transistors today; we're below the size of bacteria. We're approaching the size of a virus. In fact, at 90 nanometers, you can put 100,000 transistors on a small grain of sand. That's a lot of capability.

With the microprocessor, we built integrated circuits by putting a few transistors on a chip. Then we said, "This looks like a good building block; let's make a bunch of these." We had these little cookbooks for building circuits. It was just like having Lego blocks. You could build amazing things with those blocks. After ten years of Moore's law progress, we could all of a sudden build this thing called a microprocessor. What that let us do was bring the computer's programming environment to systems design, which did two great things. First, it raised the level of productivity of the engineers. Engineers

were now designing with these blocks that had a lot more in them. Second, we increased the pool of designers. No longer did you have to know how to build things in logic; you could be a programmer and design systems. So we increased the productivity of engineers and the pool of people that could participate and design, making the whole industry more productive. A great thing, the microprocessor. It was the second generic component. We had the transistor, then the microprocessor. Programming became problem solving. Before the microprocessor, you had to understand what structure supported the system and know how to make it run the algorithm. Once the microprocessor came along, all of a sudden someone's handing you resources that you need to do the job, and all you have to do is write a program that makes it go through the behavior that mimics what you want it to do. That grew from almost nothing, when the microprocessor was commercially *introduced* in 1971, and I'm not saying *invented* because that's a whole other story, to today, with billions of units per year in shipments. In 1999 or some time close to it, manufacturers were shipping more microprocessors every year than there were living people on the planet. Now we're probably approaching one and a half microprocessors per person per year. Pretty soon it's going to be two microprocessors for every living person on the planet, every year. That's not just the installed base. That's how many actually go out every year.

We're building chips with transistors that are so small you can't even see them with an optical microscope. We're below the size of bacteria.

The other thing that the microprocessor did for us was stall progress in design methods. For thirty years now, people have thought of problem solving as programming a microprocessor. Universities teach it; the big companies that sell microprocessors are all for it; the installed base does that; the engineers do that. We've warped the way we solve problems to suit the microprocessor. That was fine for a while. The microprocessor didn't start out in a PC; it didn't start out as the engine in a computer. It started out displacing these Lego building blocks. It started out in embedded systems. It was a low-cost, adequate performance thing. But after another ten years of Moore's law progress, a microprocessor was powerful enough to power a computer system. Then along came the PC. IBM intro-

duced the IBM PC in 1981, which legitimized the use of personal computers. It's dominated the industry for twenty years. Something like 30 percent of all semiconductor components sold in the business got into PCs. It dominates the industry.

The supply of performance grows with Moore's law. The difference between supply and demand is really important. Supply grows at some rate that in this case is determined by how fast the industry sets the treadmill. At Moore's law rate, performance is increasing. How does demand increase? Nobody measures it. Nobody really knows. But it's not necessarily tied to Moore's law. It's probably growing more slowly, but in some industries it may grow faster.

Diverging growth in supply and demand led to the value PC. When the PC came out, it just wasn't good enough to satisfy anyone.

Diverging growth in supply and demand led to the value PC. When the PC came out, it just wasn't good enough to satisfy anyone but the performance of the PC grew at some rate that was similar to Moore's law. It doubled every eighteen months. Demand was growing at some other rate. People wanted ten times what they could get when it first came out. But over time the PC got better and demand grew at some other rate. What did the administrative assistant want? What did the personnel department need? Demand spread out and supply came up faster than demand was growing. Everybody bought leading-edge PCs as long as no PC was good enough. But over time a whole bunch of customers all of sudden said, "Gee, that PC is probably pretty good. Maybe I don't need to pay premium prices for a leading-edge PC if I can get a PC that's good enough for a great price." That's the value PC; it's good enough performance at a great price.

After forty years of progress, I think the transistor is good enough. You guys are probably ready to believe me that the PC is good enough because many of you are probably satisfied with the performance you're getting from your PC and would probably rather have a value PC than a leading-edge PC. But what's this stuff about transistors being good enough? Who would ever say that? That's really stupid. Nevertheless, let's take a little closer look at it, and maybe I can convince you that transistors, if they aren't good enough today, which I think they are, will be good enough pretty soon. In fact, we may have overshot it. What's happening is, as you come down, there are smaller and smaller transistors until you get to the 90-nanometer transistor. Let's assume it costs the same to build a factory; it costs the same to build masked sets; it costs the same to do process development, no matter how big the transistors are. That's an unfair assumption, but it leads to the general case for Moore's law. It says that as those transistors get smaller, the chips get cheaper. But if you only build one transistor, it's going to cost you a lot. For one chip it's going to be expensive because the process development, the factory, the masked set are all lumped into that one chip. If you build 10-million chips, they get a lot cheaper. The problem is this isn't a good model because it costs me about twice as much each time I go down one process generation to build a factory that can do that. The equipment and the factory double with each new generation of process. The second thing that happens is I've got to build this masked set to do a particular chip. Those

masked sets do a little better than double with each transition. The third thing that enters into this is that there's a set of design rules that goes with these transistors. It shows how I build the transistors. There's just an enormously complex set of rules that goes with each transistor and is unique to each. Also the cost of development goes up each time. That cost, incidentally, is about $500 to $600 million to write the set of rules for how you build these transistors versus how you build those transistors. So throwing that into the equation, all of a sudden it's cheaper to build larger 500-nanometer transistors. And, if you're only going to build a thousand chips, those are the cheapest transistors you can get. That's because the masked costs are less; the foundries are less; the process development is cheaper; but the transistors are bigger so I have to build more wafers.

Maybe I can convince you that transistors, if they aren't good enough today, which I think they are, will be good enough pretty soon.

As I go down to 350 nanometers, I get a crossover at some volume point that says if I build more than X-number of chips, it's cheaper to build smaller transistors. With each process generation it gets cheaper. As size shrinks, I have to build more of them in order to make it the most economic thing to do. But when I go from 500 to 350 nanometers, I get this huge kick in how much I save per chip. As I move down though, I don't get as much of a kick—only a couple of dollars for the move from 130 nanometers to 90 nanometers. I save just $2 a chip instead of $70 a chip or $35 a chip.

We think about the PC market as being *the* market, but PCs are only about 150-million microprocessors a year. Where do the rest of those 8 billion chips go? That's what's driving the industry, not the 150 million that give Intel all its revenue. Two things are happening as we get down into smaller and smaller transistors. First of all, the gains are shrinking. The incentive to move in

dollars per chip set is shrinking. The second thing that's happening is that the transition points move. While I only have to build 2,000 chips in order to make it cheaper to build 350 nanometers and 500 nanometers, once I get down here to 250 nanometers versus 130, I'd have to build 64,000 chips. It's a non-linear scale. If you don't need a few million chips out of the process, then there's no incentive to go from 130 nanometers to 90 nanometers, particularly since you're only going to save a dollar on that chip. For a lot of applications the transistors are already good enough or are getting good enough.

Here's another way to look at it. There are two components to building a chip. One is how much it costs to run a wafer through the plant. That's almost fixed; it goes up a little as the process gets smaller. Now this doesn't square with transistors getting cheaper as they get smaller because fixed costs are escalating. Up until now I've been talking about it as if we built all these plants at the same time, but that's not what we did. We built the 500-nanometer plants eight or nine years ago, and it turns out they've amortized many of those costs. So for the last three processes we're still paying for that plant and equipment and process development, but for anything older than that we've already paid all those costs. You look at a chart, and say, "Gee, they do kind of get cheaper as the transistors get smaller." But once those old factories are paid for, all of a sudden 250-nanometer transistors are the cheapest transistors you can get. That's a value transistor. There are a lot of applications for it, if you're not performance limited and most applications aren't. What you care about is the cheapest transistor you can get. So there's the value transistor for 2003. When we go to a 65-nanometer process, we're going to put another value transistor out there. Just like the value PCs, the value transistor next year is going to be a better performer than the one this year. Its position is that it's the best value for your dollar.

The incentive to move in dollars per chip set is shrinking.

Microprocessors just won't do it. The reason microprocessors won't do it, essentially, is that they are simulating circuit behavior. They're not implementing things for you. To illustrate that, PC's memory is unsuitable; flash memory wears out over time; it's incredibly slow. SRAM is incredibly power-hungry and area-hungry; it forgets what it had in it when you turn it off. DRAM is volatile. But, those

are the things that locked up niches in the PC. They are the reason we couldn't get to a better memory cell because in the PC they were so good at what they specifically did that they couldn't be displaced by something better because nothing better could beat the cost performance for what they did, where they were. However, that's all changing now. What we really need in a memory component is something non-volatile, something that's as fast as SRAM and as dense as DRAM, and there are several candidates for that. There are some big backers for it. It's been around since 1988 or so, maybe even earlier, but it couldn't get traction because a PC owned the market.

Why isn't a microprocessor any good? When you're on vacation and want to take a postcard-esque picture, you can pull out your handy lens filter and camera and take a picture. Unfortunately, you probably forgot that lens filter, so when you go home and load that 4-megapixel image into Photoshop to process it a pixel at a time, it looks at a little window of 9 pixels and tries to figure out what it's supposed to look like. It runs tens of thousands of instructions per pixel on that 4-megapixel image. That illustrates for you how inefficient the computer is. That lens filter could have processed the entire wavefront in real time for no cost; there's no battery in that thing. It's just a wavefront that gets transformed as it goes through the lens. Physicists like Carver Mead can tell you what the equations of the transformation of that lens are.

Well if we know the equations of transformation of that lens, why don't we just implement them directly in some kind of hardware? That's what programmable logic does. It sits between the extreme inefficiency of the microprocessor, which is the interpretation and simulation of those equations, and the lens that actually does the transformation in physics by implementing the equations directly.

So why not just build an ASIC to do that? They are inflexible and too expensive to design. Their high fixed costs require large production runs. There are problems with programmable logic devices (PLDs) as well. They're too expensive and too slow; configuration memory is too large; it's too big; it's volatile; and the chips are dominated by wires. With all those problems why do I think that PLDs are the coming thing? If you design an ASIC, you've got those high fixed costs, but if you buy a PLD you just buy a component. PLDs are more expensive. But over time—as component costs are falling for PLDs—ASIC fixed costs, mask costs and process costs are rising. In 1997, there was the crossover when PLDs became a good design. Now, it is always better to design with PLDs. There are a lot of reasons why that is true, but I'm running short on time.

Just like the value PCs, the value transistor next year is going to be a better performer than the one this year. Its position is that it's the best value for the dollar.

ASIC capability grows at some rate related to Moore's law. But the demand for performance and capacity in an ASIC in a system is not necessarily growing at a Moore's law rate. What's funny about this is that when you go to conferences you always hear these big confrontations of PLDs versus ASICs. And they're always fought by the guys at the leading edge of demand. So it always sounds like PLDs are a stupid thing to do, but in fact most of the market is not at the leading edge. It's in some gray middle.

Mostly a lot of people are satisfied with what they get in PLDs. In fact, it's the right thing to do.

In a nutshell PCs are good enough. The emergence of the value PC shifts design emphasis from the PC to untethered systems. Because margins are going down in PCs, I've got to deploy those resources where they'll get more return. SRAM, DRAM, FLASH are unsuitable for untethered systems. A new non-volatile memory device will emerge, and non-volatile memory is going to improve PLDs. Transistors are good enough. The value transistor favors foundries over integrated device manufacturers. The fabs are good enough. Semiconductor processing equipment for a long time has been a leading indicator of the health of the industry. That's no longer true. If the fabs are good enough, the guys that sell refurbished equipment are the ones you want to watch, not the leading-edge guys that are selling to the guys that are circling the drain or have gone too far. That's not the right thing to watch. 3D wafer stacking, another whole topic by itself, will also improve PLDs.

Untethered systems are changing the design goal from cost performance to cost performance per watt. Microprocessors and DSPs just can't do it. In untethered systems demand is a stair-step function, moving up faster than Moore's law. Microprocessors are stuck on the Moore's law improvement curve. They can't do the job. PLDs and ASICs are also competing for the same designers. It's a slam-dunk for the PLDs to take over the ASIC market, which incidentally is about ten times the size of the current market. So there's a 10x market the PLDs have as a slam-dunk. The next move they make is to displace DSPs and microprocessors where they're not good enough to do the job. That's another market that's more than ten times as big as today's PLD market. The problem there is that the designers in the microprocessor market are programmers. The designers in the PLD market are designers. That isn't compatible so you need some kind of software that will bring those designers that are programmers forward to be able to use PLDs. And finally reconfigurable systems will emerge.

In untethered systems demand is a stair-step function, moving up faster than Moore's law.

Why can't we eat outside again today?

I'm confused. Is it striped-shirt day or baseball-cap day?

What the hell is this kid talking about?

Breakfast, day-two. Still saving a seat for George ...

Has anyone seen him?
He's about this tall.

Programmable Logic

positioned for explosive growth

tuesday

august 26

9:45 am

Robert Blake
Vice President,
Product Planning,
Altera Corporation

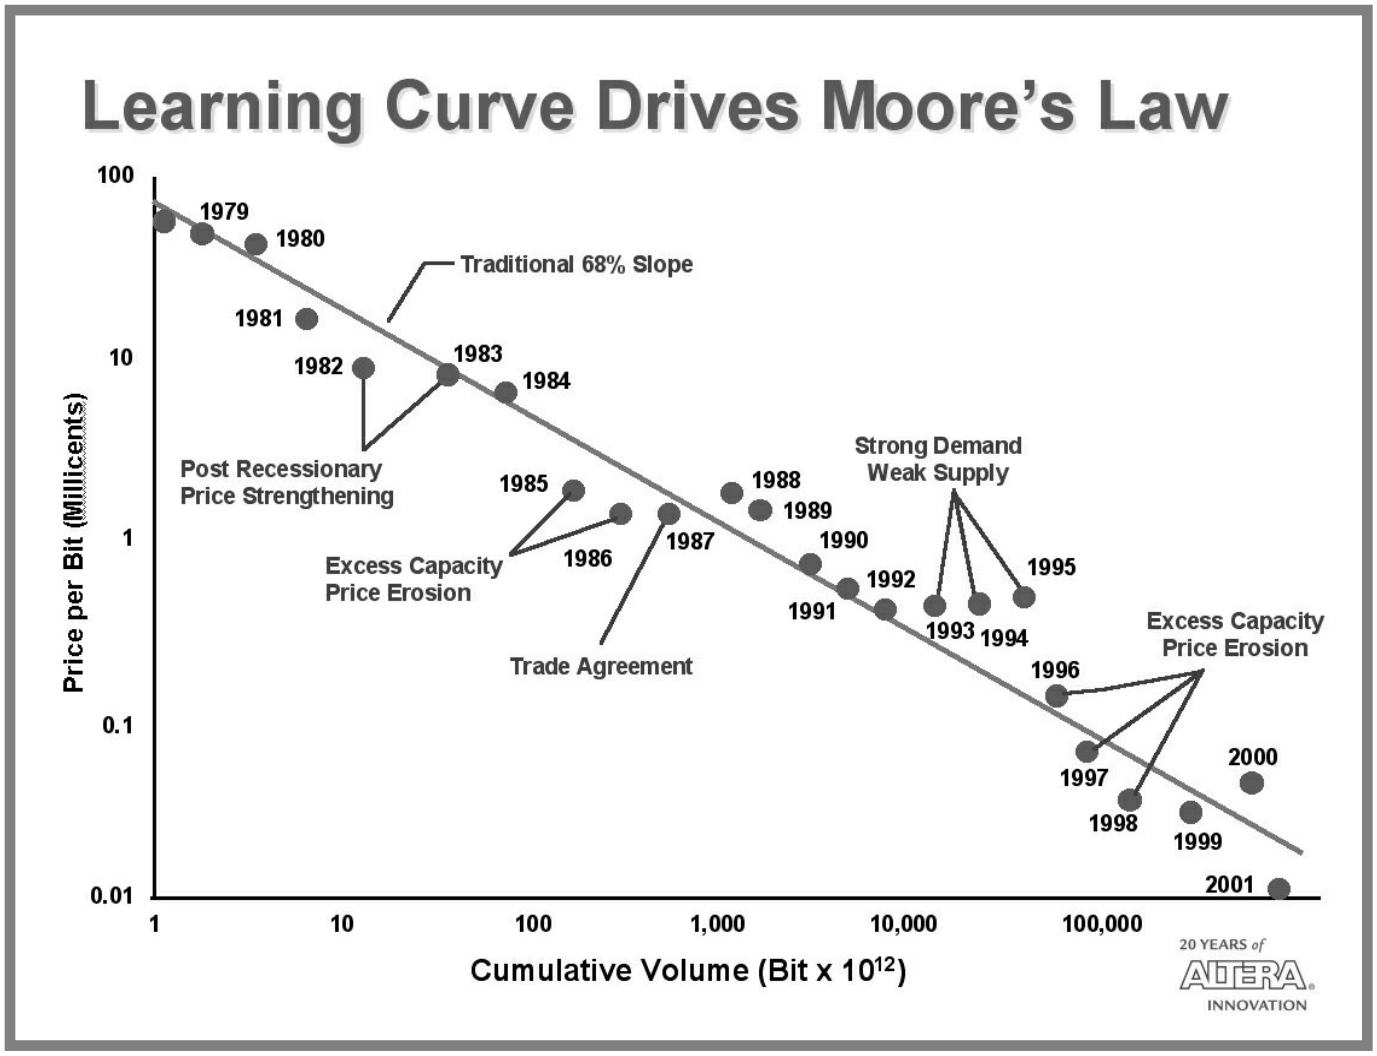

tuesday

august 26

10:00 am

Jordan Plofsky:*

My presentation is a little bit like Nick Tredennick's presentation in that I believe there is an explosive opportunity for programmable logic, although I don't believe Moore's law is a detriment to most companies. In fact, I think there are technologies that are going to have a Moore's law advantage and there are technologies that Moore's law will hurt. There will be *friends* of Moore's law and *foes* of Moore's law, and programmable logic fits into that friends category. I hope by the end of this presentation I will be able to demonstrate that to you.

First we'll look at some cost trends because that's really what makes you a friend or foe of Moore's law—whether you can afford the costs of going down the path of smaller and smaller process geometries. Then we'll look at the impact of that trend on end markets, where the rubber hits the road. I use market demand, supply, and capacity to express Moore's law. Where there's strong demand and weak supply, we get more pricing power, and we go off of Moore's law in terms of the cost per transistor and how fast that traditional slope of costs goes down. On the vertical axis, I used cents per bit on a DRAM or milli-cents per bit in the 2001 time frame, when the industry was down really low.

I don't want anyone to walk away from here thinking that Altera believes that Moore's law is not going to continue. We know it's going to continue. We have lots of technologies in the fabs to be able to do this. There have been many prognosticators in the past that have looked at the market, and at some point you can't go any further because you would draw a transistor with light that was larger than the transistor you were drawing. Now we're in the sub-wavelength lithography area, and we know we can get this all the way down to 45 and possibly below that. This is not about Moore's law stopping or process geometries continuing to go further and faster. It's about digital standard products and ASICs and a business model built around it, and those are becoming foes of Moore's law as we go forward.

The cost of coming up with a new chip, the design cost and the mask cost

* Senior Vice President, Altera

of those first few wafers, is around $30 million. Many look at this $30 million number and just don't believe it. I want to take you into the depths and convince you that $30 million is actually an aggressive best-on-the-planet number in regard to those costs. The costs are higher than that in the 90-nanometer realm. Now this is to develop a single mask set, a single circuit for $30 million. We did a cost simulation with real data. We used an 8 x 8 die because an 8 x 8 millimeter die in a fab is one of the best-sized dies to build; it's the most cost-effective. As you get larger on those die sizes, you start to get defect density yield decreases, which means you should try to get to a smaller die size, or else put up with the defect density yield decreases that exist. So I took that 8 x 8 die and I filled it up with just logic gates. If you think in transistors instead of gates, multiply these numbers by four to give you transistors. In 90 nanometers that's about the size of a Pentium IV. So at .35 microns or 350 nanometers, you could find about a half-million gates on an 8 x 8 die. At 90 nanometers you can put 10-million gates in that 8 x 8 die. We'll use that to normalize costs across process geometries as we move forward.

In this simulation of costs over time, I used very conservative, fully loaded development costs: salary; benefits; occupancy; rent; compute power; and the software to run it to design those circuits. For some customers and suppliers, these numbers are probably close to $100,000 a year. In 1997, engineers were designing in .35; in 2003 they were designing in 90 nanometers. The cost of that engineer, very conservatively, is $200,000, which will make these numbers very aggressive as we go forward.

When I was in the design groups at LSI Logic and Analog Devices, we were able to get about 30,000 gates of logic created by a designer in a quarter. If you just do the math on software of that high-level descriptive language—how many lines of code a software engineer or an IC designer can do—you're limited to approximately 50,000 gates per quarter. We've taken that number and multiplied it by five. So we're assuming there's lots more productivity that comes out of the electronic design automation industry and that there's design reuse as well. Then the design has to be verified and that's the toughest job in our business. Do you have a bug that's going to cause a recall of all those systems out in the field that's latent in your design? We used a very conservative number of 2x, the verification cost to the creation cost, and on most programs you see 4x or 5x. Then we looked at the physical design, taking it from the behavioral abstraction in software or a high-design-level language down to how you draw those transistors and the pattern generation masks to build those wafers. If you look at 350 nanometers at a half-million gates it takes 14 man-years of engineering. And at 90 nanometers it takes 52 man-years' worth of engineering. Again, we used $200,000 so that the engineering cost is $10 million to get us from just what we know we have to build to a physical design and then to mask and wafers. What you've been hearing in the press is that mask sets are going up exponentially fast. But mask sets are actually about 10 percent of the cost here. It's not

Before you spend all this money, did you do the right thing? Do you have a bug that's going to cause a recall of all those systems out in the field that's latent in regard to your design?

about the mask sets; it's about the engineering, to be able to get designs with that many transistors to work and to be verified.

If you're going to go ahead and do a digital standard product such as a set-top box decoder or a DVD standard product, it's a lot more expensive than what I assumed here. All I showed you was the logic design cost to get to the mask sets—that $10 million. But before you know it, you're at $30 million. Some of the venture capitalists are saying it's $40 to $50 million. I hope we're all on the same page, understanding that you're not going to get out alive trying to do a 90-nanometer design for a standard product or an ASIC at $30 million. If you're going to run that business at 20 percent R&D, which is on the high side of where we want to be in business models in the semiconductor business, you have to generate $150 million worth of revenue from that design for it to make sense. We don't know of many designs out there—in fact we have trouble finding any design out there—that gets more than 10 percent market share. I know there are companies that have more than 10 percent market share, but there are a number of different products that get them greater than 10 percent market share. Even the highest volume products don't get 10 percent share out of a single design, which says that you have to go ahead and target a market that's a billion and a half dollars.

I hope we're all on the same page understanding that you're not going to get out alive trying to do a 90-nanometer design for a standard product or an ASIC at $30 million.

I wanted to cheer everyone up in regard to a $1.5 billion dollar market. If I'm going to go solid chip at $10 that means we need 150-million units for that market to make sense. If it's just a 1.5-million market, which is a big market, then each chip needs to be at $1,000. Development costs either make you a friend or a foe of going down further process geometries, and it really means products must span many markets in order to get that $150 million worth of revenue out of the mask set. Either you have to have very large markets or a lot of markets that cumulate lots of units.

Based on projected 2003 data from Gartner and SIA, the communications market for all semiconductors will be about $30 billion. Now if we use that $1.5 billion criterion to say we need to invest in that market to get our

return, there are only a couple of segments of the communications market that make sense. Take digital cellular handsets at 400-million units; the total semiconductor content in a single cellular handset is $41 per unit—$6 worth of digital logic that's in that cellular handset. There are lots of communications markets where it just doesn't make sense any more to drive toward a standard product or an ASIC because the market is not large enough. You won't get your return on investment as a manufacturer. You may ask the question, "Why were all those standard products and ASICs being done before?" If you look at the normalized curves at 350 nanmeters (.35 microns), the cost to do that was about $2 million. If you want to get a 15 percent return over a five-year period, you need to generate $3 million worth of margin; it's a $6 million product. It was a much easier hurdle to get over. As we get into areas where it's multiple tens of millions of dollars, maybe even $50 million, there are less and less of these markets.

If we normalize those costs taking them from the 350-nanometer days in 1997 to the 90-nanometer days in 2003—from that $2 million to that $10 million against the growth rate of the industry—you can then see a train wreck is happening. The market is not growing fast enough and development costs are going up, causing new design methodologies to be taken into effect. You can't use the same old hammer that you were using in the past to fix this problem. It ends up showing up in earnings per share. Either you have to try to get more money out of your customers or you have to try for fewer costs from other parts of your business to support those dollars. In networking it's not as bad. We'll still see standard products and ASICs continue to be successful in networking because the gap isn't as large. In the consumer market there are a few markets that make sense for digital standard products to target, but that gap is widening. There will be more and more markets that programmable-type solutions can

Those that are friends of Moore's law will survive in the next five to ten years and grow their businesses.

attack. Industrial and automotive markets, for example.

I hope I've convinced you of what I came here to tell you; there are really two types of semiconductors. Those that are friends of Moore's law will survive in the next five to ten years and grow their businesses: programmable and analog. The programmable technologies we believe will be the winners going forward are processors because they can afford that development cost and fit into a lot of different market segments; memory, which is a programmable technology that fits into almost every market segment; and programmable logic. There's lots of potential for explosive growth in the programmable logic area because we can take that single mask set and those design costs and spread them over the customers in the various markets we serve to get that return on investment and continue down that path of lower costs, more performance, and higher density that's all the goodness of Moore's law. We will then be able to spread those dollars as we go forward.

We found George.
He was hiding
under the table.

The ballroom was crowded so we decided to sit out here.

Still stuck in the back row ...

Susan and A. Tappen Soper call home and tell the kids they're about to double their college fund.

Richard Gilder, happy he volunteered to manage the registration desk at this year's Telecosm.

The Analog Path to Semiconductor Dominance

tuesday

august 26

11:00 am

George Gilder:

Please welcome Brian Halla, the exponent of this analog revolution, and he introduced it to me and he's going to introduce it to you now.

Brian Halla:*

Thank you, George, and congratulations to you and Steve on the great turnout here, which is remarkably up from last year, and I guess that all by itself signals that we have recovered.

Last year at Telecosm we unveiled the absolute proof of when the recovery would happen. It was about this time, maybe a little later, that we predicted that it would happen two hundred and some days out. So the theme of my talk last year was "The Next Technology Boom—You Ain't Seen Nothing Yet," and during that period of time, and right up until now, we've heard all the different ways technology had died. In fact, a couple of months ago, I was with five other CEOs across the table from the Federal Reserve Board and Alan Greenspan and one of the Federal Reserve Board members leaned across the table and said, "Let's say for a minute that your industry has largely played itself out . . ." Now, obviously, it hasn't. The theme of my speech last year was that each boom gets bigger and each boom is caused by the previous bust. In fact, I said that first you have a technology boom, and we went back to the mainframe boom, the PC boom, and then the Internet boom... The fact that you have a glut, an excess, is a wonderful thing because it causes a dramatic reduction in costs, in prices. All of a sudden you have incubation of phenomenal new ideas, new ideas that drive the next technology explosion, and that's exactly where we're at now. On to the next technology boom.

Last year I actually graphed it. On the vertical was the size of the semicon-

* Chairman, President & CEO, National Semiconductor

ductor industry; we just used the semiconductor industry as a proxy for the economy and the world because it very closely aligns, even though it's only a couple percent, with the GDP. We said the mainframe boom was the first one for semiconductors because semiconductors were used to replace DRAMs in the back of mainframes or used to replace ferrocore magnets with DRAMs in the back of mainframes. When that bust happened, semiconductor technology went from being very expensive to very inexpensive, and as a result we found a way to put a whole mainframe on a desktop. We had the PC boom with 200 companies all going for 20 percent of the market. And, of course, there was the inevitable glut, the inevitable bust. But that PC boom for the semiconductor industry was about $26 billion. The fact that the prices came down all of sudden started to incubate new industries—one of those was started by two companies none of you remember, Banyan and Vines—that figured out a way to hook two PCs together to send this thing that was later to be called "e-mail" from one PC to another. And one thing led to another. There were low-cost PCs and now everyone in the company could have a PC; e-mail became a way of life. Originally DARPA had intended to tie together all the super computers to have a super-super computer. Instead, of course, that technology became the Internet and we had the dot-com boom, bust, and bubble in the year 2000. But that was wonderful. And last year we predicted that would be a cause for the next technology boom.

We all know the story. There's enough fiber in the ground already to wrap around the earth 1,600 times, and only 2.5 percent is lit. Last year I used an analogy and told everyone to read a book by Stephen Ambrose called *The Building of the Transcontinental Railroad*, nothing like it in the world. During the building of the railroad people were mortgaging their lives and everything they had to get an investment down on the building of the transcontinental railroad. All of a sudden after Promontory Point and the railroad were built, these investors in the railroad found out that they were really investing in equipment and supplies to build a railroad. After the railroad was built, there was a huge, huge economic disaster, the biggest decline in the hundred-year history of the United States, which caused this glut of boxcars, railroad tracks, and engines. The glut was a wonderful thing because as those cars went into mothballs and became very inexpensive, somebody invented the refrigerated boxcar. We had the meatpacking industry and the catalog industry started by Montgomery Ward. The period that followed immediately was the biggest economic boom in the history of the United States. It was so big, so robust that the economists invented the word "boom" to describe what was happening.

It was so big, so robust that the economists invented the word "boom" to describe what was happening.

That was the simple thesis. If you look at that equation, you should be able to put it into a formula. It's just a ramping sinus equation. So we did and came up with a formula. We have the frequency; we have the phase; we have the envelope of minimum points because we'll never go back to zero. Once we had that we simply ran it out into time to see when the next upswing would be. By the way, the first time we came up with that equation we ran it back over time to see

how accurate it was, and we were off by five or six percentage points at the very peaks of the cycles. Then Dr. Ahmad Bahai who is a professor at Stanford and Berkeley and a fellow at National Semiconductor said to me, "Brian, I've got to add chaos theory and neural nets to this to smooth out the points." Dr. Bahai, by the way, did his thesis at the University of London on chaos theory and control systems. His ideas later turned into a startup out of Texas that predicts the health of various companies and guarantees the results of investments in them. So when Dr. Bahai added these neural nets this thing became very sophisticated and we ran the data. We ended up with something that began to look compelling. We blew that up and predicted when the recovery would occur. We knew it was somewhere close to the end of the second calendar quarter. My anniversary happens to be on the twenty-first, and in thirty-four years of marriage, I've never gotten it right. I figured this is a sure way to remember it, so we picked June 21, and half of the people took it as tongue in cheek, half of the people took it as gospel. If you listen to the industry pundits today, if you look at the title of this conference [The Turnaround Telecosm] and the size of the crowd in this room, the only thing you can figure out, in fact, is that we were right. That's why we changed the title of this presentation to "Out of the Term, Onto the Next Technology Boom."

Let's talk about a few things that might make this happen. We obviously know that the mainframe, the PC, the connected PC to the Internet, and the cell phone all added to that last dot-com boom. Now we're all saying, What's next? What's the next killer app? What's going to drive the next recovery? It's pretty straightforward if you think about it.

After this conference, I'm invited to Europe to give a keynote speech at CeBIT. I knew that the transcontinental railroad was not going to go over as well in Germany as it did over here, so I couldn't use that. I tried to come up with another way to describe or to triangulate why there would be an economic explosion right around the corner. I picked a different way of explaining it. The theme of that presentation combines with my title for this year's Telecosm speech, which is "Why Sensors in Proximity Signal Processing Will Drive the Next Information Technology Explosion." The key word there is "information." We'll spend the bulk of the time at 30,000 feet talking about that, although I'll dive down to show what proximity signal processing really is.

The book I'm recommending for you this year is called *The New Renaissance*, a whole new way of predicting the next information technology revolution. It's by Douglas Robertson. The book was written in

1998, which is important, and I'll let you know at the end of the presentation why. Robertson's approach is that the only thing that limits the information technology explosion is the amount of information that we have available at the time people start innovating. In fact, he says there are four great inventions in the history of the world. The first invention was speech, around 10,000 b.c., and that was at the beginning of mankind in general. The second invention he describes was at the beginning of civilization; it was writing because writing suddenly expanded the amount of available information. The third and probably the most demonstrative invention in terms of catapulting the amount of information and subsequently the amount of innovation was the invention of the printing press, which happened to be done in Mintz, Germany. Mintz was just down the street from where I was giving my pitch so it worked out very well. And, of course, computing. Now if you look at the date for computing—1950—we haven't been at it very long. We've really only scratched the surface. There are multiple kinds of computing, but if we graph the amount of available information that currently exists, it's somewhat tongue in check because obviously it can't be accurate. When language was invented, the total amount of information on the face of the earth at that time was about a gigabyte total. When we moved to the invention of the written word, there was suddenly a dramatic expansion in the amount of available information because you could write something down and pass it around. The total amount of information became 100 gigabytes. With the invention of the printing press, the amount of information exploded because it could be duplicated and sent to all parts of the earth. And, of course, with computing we're now at a point where information is doubling worldwide every couple of years, probably even faster than that.

At the time the printing press was invented in 1450, there were about 300,000 books available on the planet.

If you get nothing else out of this presentation, think of real-time proximity signal processing.

Over the next 50 years, over 8 million books were printed. During this period of time called the Renaissance, a number of key things happened. Christopher Columbus used multiple books that were available to him in formulating his belief that the earth wasn't flat. Copernicus used multiple books that were available to him to develop his theories on astronomy. Wonderful things happened because of this availability of information. During the Renaissance, people came to realize that dragons and unicorns were not real. Now remember, the computer has only been around for fifty years. The first computers were used to do number processing; the first PCs, word processing. And now computers combined with proximity signal processors are dramatically and exponentially expanding the amount of information available to each and every one of us.

We can now safely say that connected computing will give all of us instantaneous access to whatever information is available. Image capture makes all of us with our handsets individual printing presses. Each of us becomes a printing press—every man, woman, and child who has an imaging device in their handset or carries around a digital camera—creating information every time we snap a picture, especially if we put it up on the Internet. If you get nothing else out of

this presentation, think of real-time proximity signal processing.

What's next? Information availability drives technology explosion. Information caps, or the caps on information, stifle technology explosion. Here's an excellent example in Robertson's book: if Ramses and Pericles had had the same amount of information available to them as Morse and Edison had in their time, then the

telegraph and the phonograph would have been invented back in the very early days. But Ramses and Pericles didn't have that information; they just had the raw parts. My thesis is that an infinite amount of information will require an infinite amount of processing power and an infinite amount of bandwidth, which most of you would probably agree we already have. However, it's essential that we find a *use* for the infinite amount of processing power and bandwidth already out there.

Digital is the brain and the rest of the world and the rest of the body is analog.

Creating information is the next big thing. I'm going to introduce a couple of concepts. One is not new to anyone here; it's basic sensor technology. I'm talking about signal capture, image capture, temperature sensing, and all kinds of sensing. The other idea I'm going to talk about deals with the signals right at the point at which they're sensed. So back to analog. We all know that standard linear analog consists of four parts: signal conditioning, which is amplifiers; signal conversion, As to Ds (analog to digital), Ds to As (digital to analog); signal interface, which is the serial bus or parallel bus and SERDES; and then whatever is surrounded by regulators and references of voltage, which is standard linear analog. The other part of the analog business includes ASSPs.

Digital, or what we think of as digital, consists of CPUs—central processing units—and digital signal processors. My thesis is that they co-exist. One doesn't replace the other, but one can certainly augment the other. The argument we had earlier means we don't believe that a cell phone goes on a single chip. We believe it goes on two chips. After doing the math and the investigation, we decided to spin off our GPRS chip set business. It was simply too competitive and did not add value to where we were going.

We'll now introduce a new analog capability signal capture or sensors and proximity signal processing. I use the Vitruvian man from Da Vinci to explain this concept. One of our employees said to me the other day, "So, you spun off the GPRS business and the i-businesses. Does this mean that digital is dead?" And I said, "No. Digital is the brain and the rest of the world and the rest of the body is analog." The brain is the central processing unit and the DSPs are the signals it receives come from all parts, from all extremities. What we want

to demonstrate is that the brain doesn't have to do everything. Let's give the Vitruvian Man opposable thumbs. Let's even give him intelligence right there at the fingertips.

Last year I introduced what I thought was *the* killer app. It was Given Imaging's capsule endoscopy pill. You swallow it; it shoots four pictures per second for six hours, going on to twenty-four hours; it's already on the market; it costs $450. The reason I call it the killer app is once it can do endoscopy and proctoscopy, it eliminates the alternative procedure, which at age 50 we all start adding to our physicals. It shoots pictures and this is the first generation. It has a couple of batteries; it has white LEDs; it has white LED drivers; it has a digital imaging device; it has an ASIC; it has a radio. It shoots four frames a second; those pictures are transmitted to a hard disk worn on a belt and terabytes of photographs are stored on the hard disk. There are doctors today who do nothing other than look at those pictures to try to spot anomalies. You can imagine how long it would take to look frame by frame at terabytes of data.

Let's look at a future generation of this pill, and this is new since last year. First of all, the pill takes a picture. Second of all, it starts to understand what it just saw. It looks for anomalies in the color. Certain types of tumors, such as a pancreatic tumor, release different color bile than is normal. It takes a picture and right onboard that pill will actually do color spectroscopy, checking for a color that could be out of range. If it happens to find a color out of range, the next step is to check to see if there's been any temperature variance. National Semiconductor produces a temp sensor that can detect temperature changes of a millionth of a degree. We can do onboard temperature sensing, and if we have a color variance that's out of range and a temperature that doesn't look right, we can flag that picture and send that picture and only that picture, or that set of pictures, to the hard disk for the doctor to review. Ninety percent of the diagnosis is already found inside the body, inside that pill, using proximity signal processing.

Another example I talked about last year was one of the more obvious ones. Today as we go forward, virtually everything that is subject to a terrorist attack or to any kind of destruction will have imaging devices involved. National has an application of a camera that looks both ways on the Trans-Sahara pipeline, up and down 500 meters of pipeline. Every 500 meters you have these dual cameras that do onboard image checking and visual recognition. As long as the picture looks like the pipe with nothing bothering it, nothing happens. As soon as somebody comes up to try to harm the pipe, a signal is sent up to a satellite, an operator is told to turn on a monitor, and a helicopter is sent in. This is very real. Since the last time we talked, $22 million worth of cameras have been installed on seven of the Bay Area bridges. They can virtually look at every square inch of a bridge day or night and can see a face a mile away. They can check that face against a list of known faces, and they can do that in complete darkness.

Let's give the Vitruvian Man opposable thumbs.

Image capture, image sensors, can go a long way toward protecting railroads, pipelines, bridges, and other things that could be subject to terrorist attack.

Another example is RFID. We all have RFID. National, among several other companies, is a player in RFID. But why just have a tag attached to a box

that when ignited by an interrogation beam says, "Here I am"? RFID can do a lot of other things. The example I use is an RFID tag stuck to a meat carcass that's in the back of an eighteen-wheeler coming across Iowa with a temperature sensor. Now it goes past the weigh station; the interrogation beam turns on the RFID; the RFID says that before I say, "Here I am," I'm first going to check the temperature of this slab of beef against the last several temperatures to see if that temperature is out of range or if it's dropping. If there's something happening here, at a minimum I can send a message to the truck driver telling him to give me more air conditioning. Or if I do have maximum air conditioning, I can tell the driver to pay extra to each of the weigh stations to get this truckload of beef to its destination before it starts to spoil. Proximity signal processing.

All of this will invoke concerns about privacy.

All of this will invoke concerns about privacy. A recent *IEEE Spectrum* notes the fact that, if you have a GSM handset and it's turned on, somebody, somewhere knows where you are. And the reason for that is the way your GSM packet works. Every time a cell gets access, there's a time-stamp put on it. Now typically about three cells will access a cell phone conversation, and the optimal one is selected. The other two also time-stamp the packet that arrives. That means there's a difference in the time of arrival; you can triangulate the difference with a couple of scanners and know exactly where the person with that GSM handset is. Interesting, but as [Sun CEO] Scott McNealy said, "Privacy, get over it. You lost it several years ago." As I said, every man, woman, and child with a handset, with a Foveon imaging device in it, becomes a printing press, creating information.

To go a little further into handsets: I want to look at all of the handset capabilities—our analog functions. Why can't a whole handset go on a single chip? For one reason, the analog has more and more to do in differentiating that handset. And the CPU, though it's got to crunch a lot of data, has less and less to do with the differentiation. When you get to the point of one analog chip and one CPU, if you put both on the same chip and surround the analog with noise isolation, the chip becomes prohibitively large. It becomes more expensive than the two chips in one package. We believe all analog functions belong on a single chip. The PowerWise Initiative that

National launched with ARM was pretty straightforward. The ARM Intelligent Energy Manager resides on an ARM processor today. It looks to see what the ARM processor and the handset are doing and what the repeatable, predictable sequence of instructions is because almost everything a handset does has a repeatable, predictable sequence of instructions, like an MP3 download. It knows that it has to execute those instructions in a certain period of time. It finds that it typically has time left over. In the new PowerWise Initiative, National put an implant into every ARM processor. The Intelligent Energy Manager says you have this much time to do the next sequence of instructions, but it only really takes you this long to do it. Then it sends a message to the implant, which in turn sends a message to National's power management unit that says, "Slow down the clock, slow down the frequency, clock by clock." This allows me to use the minimum amount of juice to execute that sequence of instructions. Initially, there is a savings in battery life of 25 percent. Over the stages of this relationship, you'll have 4x or 400 percent more battery life per handset.

Getting back to handsets. In display systems, everything's moving to color and better color and lighting solutions. With amplifiers we're now boosting the white LEDs to become flash attachments for handsets. We've just launched an amplifier in a microphone. We've eliminated the JFET (junction field effect transistor) inside the microphone, inside the handset, such that what comes out of the microphone is now a digital signal. It's virtually impossible now to get the same kind of noise and static that we've all complained about. And just for fun, I haven't heard a handset go off in this audience yet, but it's inevitable. Every time I hear an obnoxious ring tone—a polyphonic tone—I love it because it's another seventy cents for National.

Where are we going? You can now download your favorite MP3 clip and play five seconds of Santana when your phone rings or relive the sounds of your safari in Africa. For those of you who have walked on the Great Wall of China, you know that everybody else walking on the Great Wall has one thing in common. They're all talking on a handset, and you'll hear a ring tone that sounds something like this (plays Chinese music). There is am amp in the microphone.

There is the story of someone who used his handset not as a printing press but as a crime stopper. A fifteen-year-old kid was being abducted. Not only did he take a picture of his potential abductor, he also snapped a picture of the guy's license plate; the guy was apprehended just a few minutes later. Pretty cool. Police said that was the first time they'd ever seen anything like this.

We now have about five designs for the imaging sensor that looks out for the white line on the nation's roads. It turns out that only in the United States do we have pavement bumps that wake up drivers if they start to go over the white line. Everyone else will now have a camera looking at the white line, and if you begin to drift over it without putting any kind of pressure on the wheel, then you get simulation of going over bumps as the brakes are tapped and the wheel vibrates. The next stage, obviously, is as you start drifting over, and if you don't have pressure applied to the

Every time I hear an obnoxious ring tone—a polyphonic tone—I love it because it's another seventy cents for National.

steering wheel, the logical thing will be to bring you back into the lane and monitor your speed based on where you are, vis à vis to the car in front of you, until you wake up.

I said I'd return to Doug Robertson's book, which he wrote in 1998. I was thinking this morning as I was watching CNBC before flying up here about the UN inspectors finding trace elements of radioactivity on machinery that had been shipped into Iran. The obvious conclusion is that Iran is building nuclear weapons. And Iran denied it. In fact, they said they would give the UN snap inspections so that it would feel comfortable that there's nothing going on. I recalled that when Iraq said that they would let us do fly-overs, that that was a good first step. If you think about it, we could have put imaging sensors anywhere and everywhere, and Iraq would probably have been happy with that as well. Once a missile got loaded on a truck, then we could take action. But the interesting thing is this book was written in 1998, and the author says that with all of the information available today it's virtually inconceivable that there will ever be another war on this planet. If you think about it, this was all about lack of information. We're at a point now where information not only doubles every couple of years, but there's an explosion, an infinite amount of information being created. The next technology explosion is upon us. We've already seen the turn. It's driven by an infinite amount of information, but it has to be useful information. If I just use that pipeline example and if all those cameras were sending real images, that by itself would clog the pipeline. But they're not. They're sending images that have been pre-processed with sensor and proximity signal processing.

There's an explosion, an infinite amount of information being created.

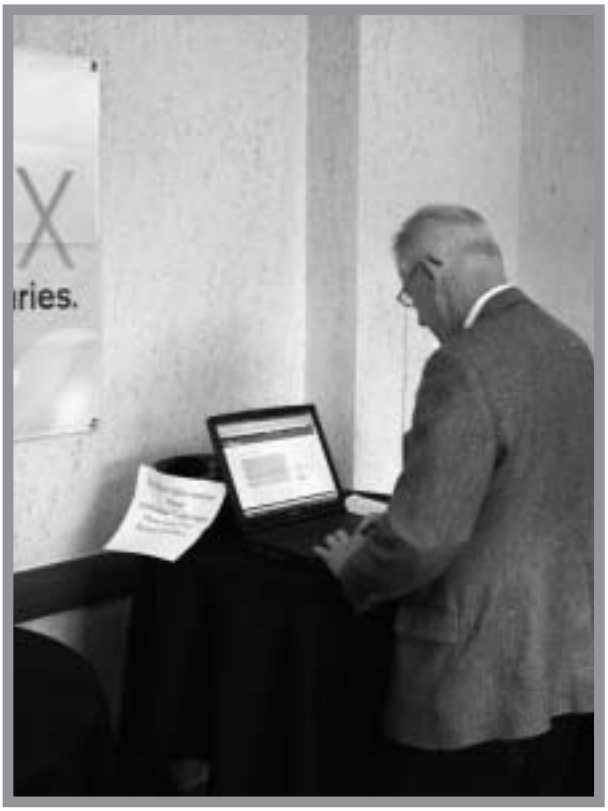

Can't anyone go a couple of hours without checking their email?

So, when are you saying I need to be on stage again?

Rich Karlgaard reviews the day's racing forms before catching the shuttle to Reno.

Analog's Ascent

analog enters the digital realm

tuesday
august 26
11:30 am

Thomas Lee:*

I think professors typically suffer from "professor's disease," which is to be a bit pedantic, so I just ask you to bear with me for two slides defining what analog is. I thought that this was going to be superfluous, but I ran into a VC in the hallway a couple of days ago who asked me to define what analog was really; it was a bit appalling because this gentleman has invested in a lot of analog startups, so I think it's a good idea to define what this is.

Analog really has to do with things that are created in nature. They're of the physical world. These are parameters that span a continuum of values like pressure, temperature, voltage, and so forth. For example, Wi-Fi signals might span a range from micro-volts at the antenna, at the limits of sensitivity, and power lines are on the order of a mega-volt. Analog spans that entire range. Analog systems also have an interesting characteristic in that they typically have very high functionality per unit of power. They are very specialized systems. An example I gave last year was the human auditory system, which manages its miraculous performance on just tens of microwatts, something that we can't even approach with silicon-based analogs today. The prob-

* Associate Professor, Stanford University • Co-founder, Matrix Semiconductor

lem is you can't really process and replicate analog signals with great fidelity. That's a difficult nature of analog systems, and because they're highly specialized it's very hard to reuse circuitry. It's hard to imagine, say, using your nose as an eyeball. That's the problem with analog, but nature is analog at its core.

Conversely, digital is unnatural. I don't mean that as a value judgment, although sometimes your laptop may seem to be satanically possessed. In digital circuits, the systems typically consume relatively high power per function, and that's an artifact of our insistence that digital systems be very flexible. Because they are flexible, we allow ourselves to employ that flexibility or exploit that flexibility. But paradoxically, that very flexibility has killed off, arguably, circuit innovation for three or so decades because we're just living off of the fat of the land, so to speak. Digital is unnatural; analog is natural. That commonly sets up a false dichotomy. You have lots of panel discussions that pit analog against digital. I'd like to offer a different view of that contest, if you will.

Analog processing of some kind, of course, is absolutely inevitable any time you want to process things that arise from the physical world. You might hear things or read things about all digital or softer radios, but those are marketing terms that really have no connection to reality. Digital processing is a powerful and incredibly flexi-

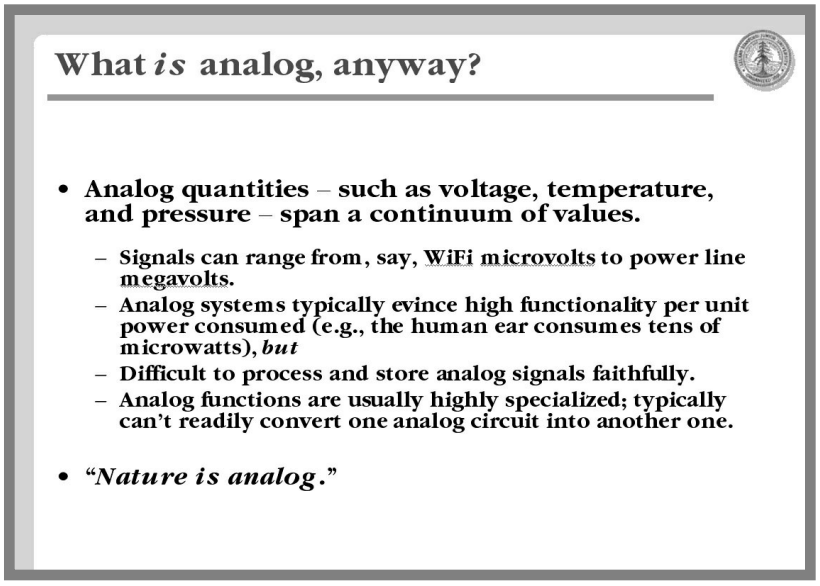

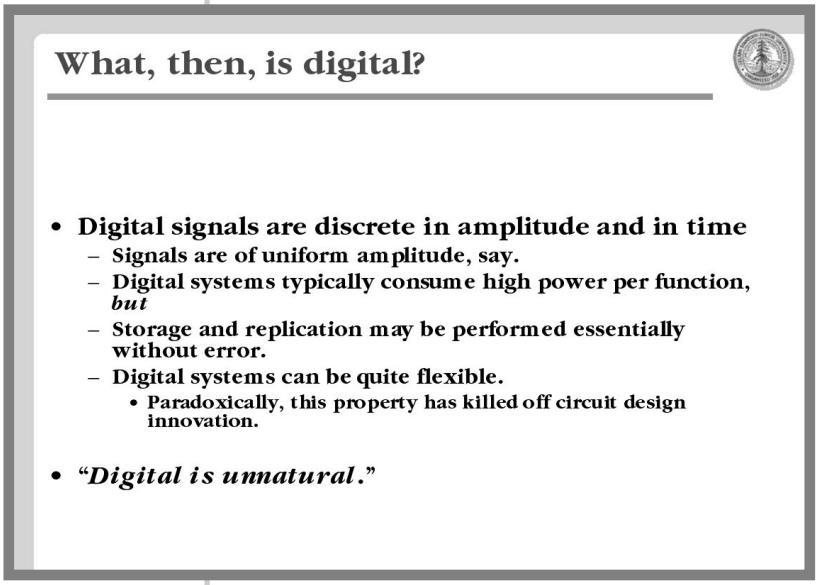

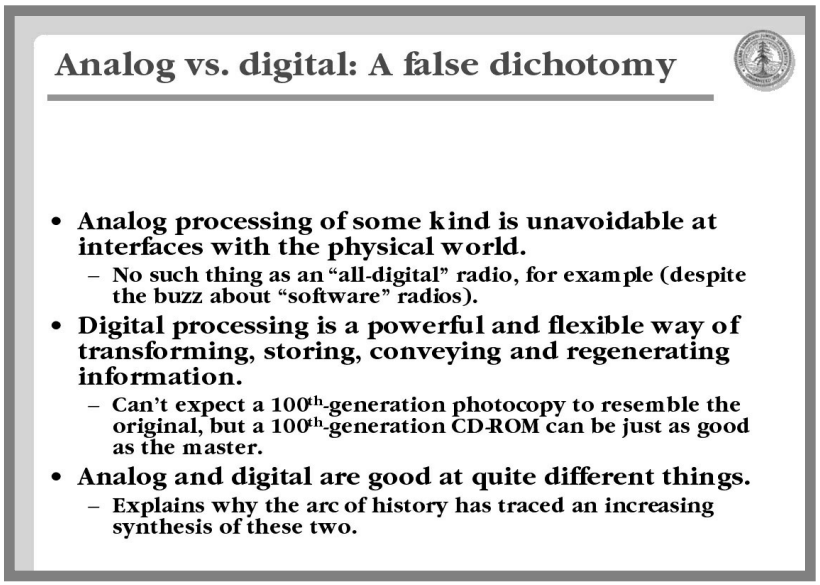

So, (why) is analog ascending?

- **Purely digital paradigms are limited.**
  - **The Pentium 4 consumes ~100W, the human brain but 25W.**

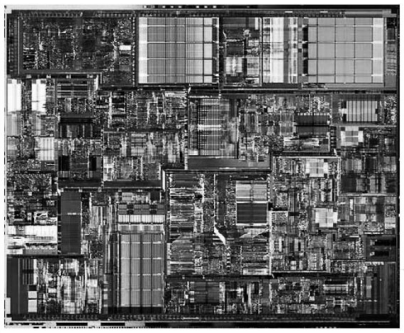

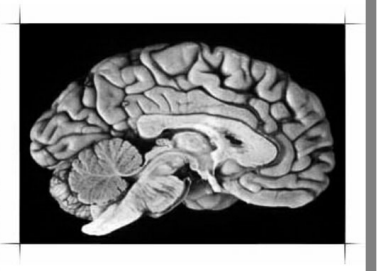

Biology has thrown down the gauntlet

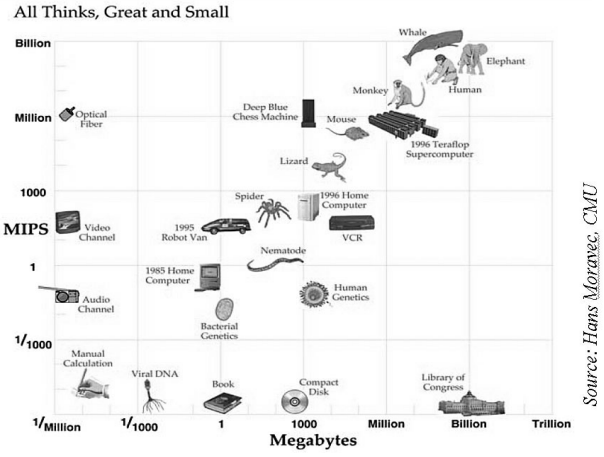

Whither analog?

- **One view: Analog will serve as merely another rationale for Intel to build ever-more powerful CPUs. This co-dependency will continue until market forces dictate otherwise.**
- **Another view: Insights gained from studying diverse fields will inform, and drive, a convergent evolution of seemingly disparate technologies.**
  - **Hard to predict where this will go, but my guess is that ethicists and social scientists will have plenty to chew on.**

ble way to transform, store, convey, and regenerate information. Just as a very simple example, a hundredth generation Xerox copy bears no relationship whatsoever to the master, but you can make a hundredth generation CD-ROM that is identical to the first-generation copy. So analog and digital are good at quite different operations and to insist that I only do things in analog form or only do things in digital form is to preclude a whole collection of very useful behaviors. An example to show how the arc of history has, in fact, recognized this truth is Bell's prototype phone evolving into today's cell phone with a color video screen. If you look at the guts of a cell phone, only a small percentage of the silicon is in analog form, but it's a critical piece of the puzzle. We couldn't possibly have a pure analog cell phone that does all the things that the modern cell phone does, nor could we have it be all-digital. You need the synthesis of the two technologies.

Another example in wireless is the Wi-Fi cards that [Intel COO] Paul Otellini mentioned earlier. Wi-Fi cards are commoditized before they're productized. They are $19,95 at Frye's, and I don't know who's making money on this, but as a consumer I'm happy that that's all I have to pay. Another example is the gramophone, which was replaced by today's CD-ROM. Without digital we couldn't possibly have these things. It's a synthesis.

Because analog resides at that critical interface between the physical world and

I thought someone said there'd be snow. Why'd I bother stuffing a ski boot in each bag?

digital processing systems, as [National Semiconductor CEO] Brian Halla mentioned, sensors are a key part of the story. We're able to leverage the tremendous investment in digital process technology to enable new things in the analog realm. As one example, look at an array of ultrasonic transducers. You might ask, "What good is that for other than publishing papers about it?" Well, you can do things like transform an early-generation ultrasound, where you can barely discern the presence of a fetus enough to say, "Yes, you have a fetus" to an ultrasound using modern sensors coupled with powerful digital signal processing, so not only can you say, "Yes, there's a fetus," you can discern its gender, and if you look very closely, even its physical characteristics.

Is analog ascending and if it is, why? A paradigm limited to either analog or digital alone is limited in its power. A Pentium IV consumes something like one hundred watts of power. A human brain consumes about twenty-five watts; some people less than twenty-five, a function of how close you are to the Beltway, perhaps.

So analog and digital things are good at quite different operations and to insist that I only do things in analog form or only do things in digital form is to preclude a whole collection of very useful behaviors.

The Pentium consumes four times as much power, and yet it has never created anything as close to the exquisite transcendence of say Brahms's *Trio No. 1 in B major Opus 8*.

If you allow yourself to think so far outside of the box that you can't even see the box, you might plot a chart with some person's estimate of the number of instructions per second required to mimic the behavior of those various organisms and mechanical objects on the vertical axis, and an estimate of how much information storage there is in that object on the horizontal axis. In the upper right-hand corner you would have whales, elephants, and humans, but not too far down to the lower left would be systems we're able to build today. There are three orders of magnitude spanning each major division in the scale, but it indicates that in the not too distant future, within a decade or two, we will have in our grasp the ability to produce computations that are of an order of complexity that we estimate to be similar to that of humans and high-order mammals.

Does that become a future business? Does this become something we have to worry about? I'm allowed to think crazy because I'm a faculty member and don't have to do anything useful. My students will back that up.

One view is that analog will continue to serve as simply another source of bits to be processed by CPUs made by Intel. We'll all enrich Intel and the rest of us will just barely survive. Another view is that insights that are being gained from studying multiple paradigms and multiple dimensions, biology, genetics and so forth, will inform and drive a convergence of disparate technologies to who knows what future. We don't know. I can predict this, though lots and lots of papers and very, very well exercised discussions will take place among philosophers, ethicists, and social scientists. From a business point of view, I still wonder why no one has addressed the following observation made by rocket scientist Wernher von Braun several decades ago. "Man is the only computer that can be mass-produced by unskilled labor." You have to compete against that truth.

'Man is the only computer that can be mass-produced by unskilled labor.'

I love Tetris .

This would be a lot more fun if I wasn't supposed to be working.

Why are you standing over there? I saved a seat for you right next to me.

Once again, Tredennick forgot to pack his dark socks.

I told you I wanted to sit on this side of the Altera guys.

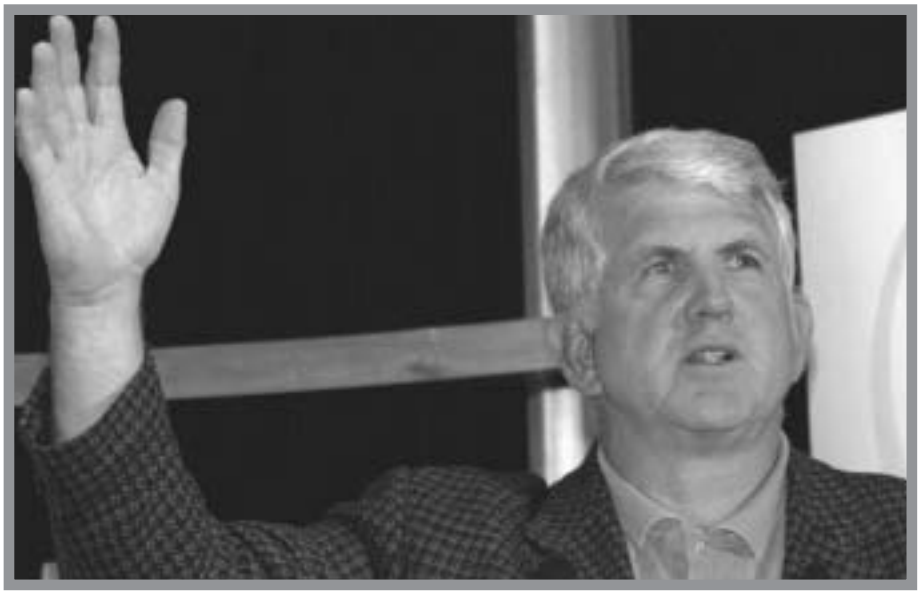

Lunch with Bob Metcalfe

what the past thirty years of ethernet says about the next thirty years of . . . life on earth?

tuesday
august 26
1:00 pm

Rich Karlgaard:*

What the past thirty years of Ethernet says about the next thirty years of . . . life on earth?

Bob Metcalfe:**

Ethernet, the networking standard, was invented in a memo I wrote on May 22, 1973, at Xerox Park, that would be thirty some years ago. George heard about this so he invited me to come here and talk about what lessons I might have learned over those thirty years, and then I tacked on to the end of my title, "Lessons for the Next Years Here on Earth." I added the phrase "here on earth" because I only have thirty minutes, and I wanted to focus my comments. So I'm going to give a little history of Ethernet and talk about the ensemble of ideas, winning ideas, from which it has benefited. I'm going to talk about the business model of Ethernet; about the silver bullet, which I call "FOCA-CA"; then I'm going to take a shot at monopolies, just for your enjoyment, George; and finally I'm going to kick some universities in the shins and wrap up right there. My goal is to be done by 2:15 p.m.

As Rich [Karlgaard] pointed out, Ethernet, the CSMA/CD local area network, IEEE 802, and the Internet and TCP/IP wide area network were both invented in Palo Alto in 1973. Neither was invented by Al Gore, although I have to admit he later contributed some very clever "AlGor-ithms."

I invented Ethernet on May 22, 1973. That sentence is a factoid. Do you know what a factoid is? Anything that has been said three times. I not only invented Ethernet, I invented that factoid. It was in the mid-eighties and I needed a hook for some particular event, and the hook was to be Ethernet's birthday so I came up with the idea that Ethernet was invented by me on a single day, May 22, 1973. But clearly, things like Ethernet are never invented by a single person on

* Publisher, *Forbes* magazine
** Internet Pioneer • General Partner, Polaris Ventures

a single day. That's just a marketing thing, and we've been celebrating it every five years ever since then. Ethernet wasn't exactly invented by me; it wasn't exactly invented on May 22, 1973; and actually the Ethernet we have today isn't anything like the Ethernet that was invented in 1973, but the factoid stands.

Now, let me give you the real history. I'll touch on just six emergent points. The Ethernet, a networking standard connecting computers within buildings, was developed from a set of ideas borrowed from the ARPANET, which you may recall is the early version of the Internet, and the ALOHA Network, a packet-radio network circa 1970 at the University of Hawaii. I spent a month in Hawaii studying the ALOHA Network (I recommend field trips like that). Dave Boggs and I built the first Ethernet starting in 1973; then Gordon Bell, who was then the vice president of DEC, began standardizing Ethernet with me in 1979; and next Ron Crane—who helped me start 3Com—and I shipped the first big-time Ethernet product, the EtherLink, in September 1982. EtherLink was the first Ethernet for any personal computer. We had been shipping tens of units and then we suddenly started shipping thousands of units and a couple of years later millions of units per month. By 1983, there were people buying Ethernet whom I did not know. By 1986, there were people inventing Ethernet whom I did not know. Overall, I sold my first Ethernet adapter in 1981 for $5,000; it was made of TTL, by the way, since you asked. Today Ethernet adapters go for one-hundredth to a thousandth of that price and last year 100 million—actually more than that, but we only deal with round numbers—new Ethernet connections were sold, representing more than $10 billion in revenue for just the ports alone. And that's great news, but I did not get the usual 10 percent cut of that $10 billion.

People have asked me how much money I've made on Ethernet so I'm going to answer that question right now.

Unlike Moore's law, Metcalfe's law has never actually been true.

What did I get for inventing Ethernet? Xerox paid me $250,000 total, salaries and benefits, for the eight years that I worked there on Ethernet, and I got stock and salary and benefits and options in 3Com Corporation—that was pretty good—and at my age now I get various lifetime achievement awards, which are now my duty to collect. But most of what I got, of course, was having the honor of having a law named after me by George Gilder. And while we're on the topic of the law, I hasten to add, especially for those of you who insist on using Metcalfe's law in your slides, Metcalfe has an "e" on the end of it, and Metcalfe's law says that the value of the network grows as the *square* of the number of users. And I'm a mathematician: square is not exponential; square is polynomial. So stop this no "e" exponential thing because that's not Metcalfe's law. Metcalfe's law has another feature relative to Moore's law. Unlike Moore's law, Metcalfe's law has never actually been true, but I remain honored, George, to have you name it after me.

George Gilder:

I revised it to make it true.

Bob Metcalfe:

But now that's the Law of the Telecosm.

George Gilder:

Well, I tried to make it Metcalfe's law, but you refused to accept that.

Bob Metcalfe:

So like Ethernet, this thing going under my name now bears no relationship to what I came up with. So younger smart asses, who remind me of myself, will tell you that Ethernet bears no relationship to the CSMA/CD local area network running over a kilometer of coaxial cable at 2.94 megabits per second that Dave Boggs and I built starting in 1973. Of course, they weren't there. They didn't notice the novelty of Ethernet.

The solution was straightforward engineering. So the smart asses are right.

The novelty of Ethernet was the problem, not the solution. Never before had anyone been presented with the problem of having a computer on every desk. In 1972, there were no personal computers. Suddenly we were about to fill a building with personal computers and I got the problem. We had a laser printer and we had a time-sharing system and we had access to the Internet. My job was to build a network that would connect them all. It was a fast laser printer so it couldn't run at kilobits per second; it had to run at 2.94 megabits per second, actually. The good luck there—and there's a lot of luck in life—was to be given a novel problem. The solution was straightforward engineering. So the smart asses are right. Ethernet—10 megabit; 100 megabit; 1,000 megabit; 10,000 megabit; coax; thick and thin; twisted pair; shielded and unshielded; optical fibers; and now radio again—today's Ethernet is nothing like the one we built. By the way, Wi-Fi used to be called Wireless Ethernet, but the smart asses decided to change its name so now they call it Wi-Fi, and I'm frequently faced with the question, "Do you think Wi-Fi will kill Ethernet?" OK, so you can't say that today's Ethernet—you have it here in your hotel room, don't you?—that black thing that says "Ethernet" twice on it, *my* word—is a CSMA/CD coax network with collisions. In fact, there hasn't been a collision reported on an Ethernet in several years

now. You might say that Ethernet is synonymous with local area network, but even that's not true because Ethernet has penetrated the metro area and now even long haul as an interface. You might say that what we mean by Ethernet is the packet format that has endured for some twenty years, even to some degree on Wi-Fi—the 48-bit addresses and stuff—but I don't think that's the answer.

What I think the word Ethernet means today is an ensemble of winning ideas, and I'm now going to tell you those ideas. These ideas have allowed Ethernet over that thirty-year period to beat, kill, eviscerate Sneakernet; IBM 3270; Hyperchannel; Arc.net; OmniNet; PCNet; General Motors Token Bus; IBM Token Ring; FDDI; ATM; the list goes on. And do you want to know the next three things that these winning ideas are going to kill? Fibre Channel, which is at the basis of storage area networks; SONET, which is an old, complicated, expensive standard of the telephone monopolies—it's going down—and contrary to representation here yesterday, cell phones. Now this won't happen soon, but in the future cell phones won't be derivative of a "G" anything. They'll be derivative of "802-dot" something in the long-term future, that is, cell phones are going down. Actually, they'll be going up, but they'll be using some derivative of Ethernet, which will bear no relationship to what I invented, but they'll still call it Ethernet.

There are five winning ideas that Ethernet has used over the years to persist. Quick summary: packets; layering; decentralization; the Ether; its business model. Now Ethernet got its idea of using data packets from the ARPANET and from the ALOHA Network, and these are data packets as opposed to voice circuits or video channels. It's turned out to be a really good idea. Voice video and data will be converged when voice and video are data, which is what's happening now, and Ethernet is designed to carry them.

What I think the word Ethernet means today is an ensemble of winning ideas.

The second good idea is the layering

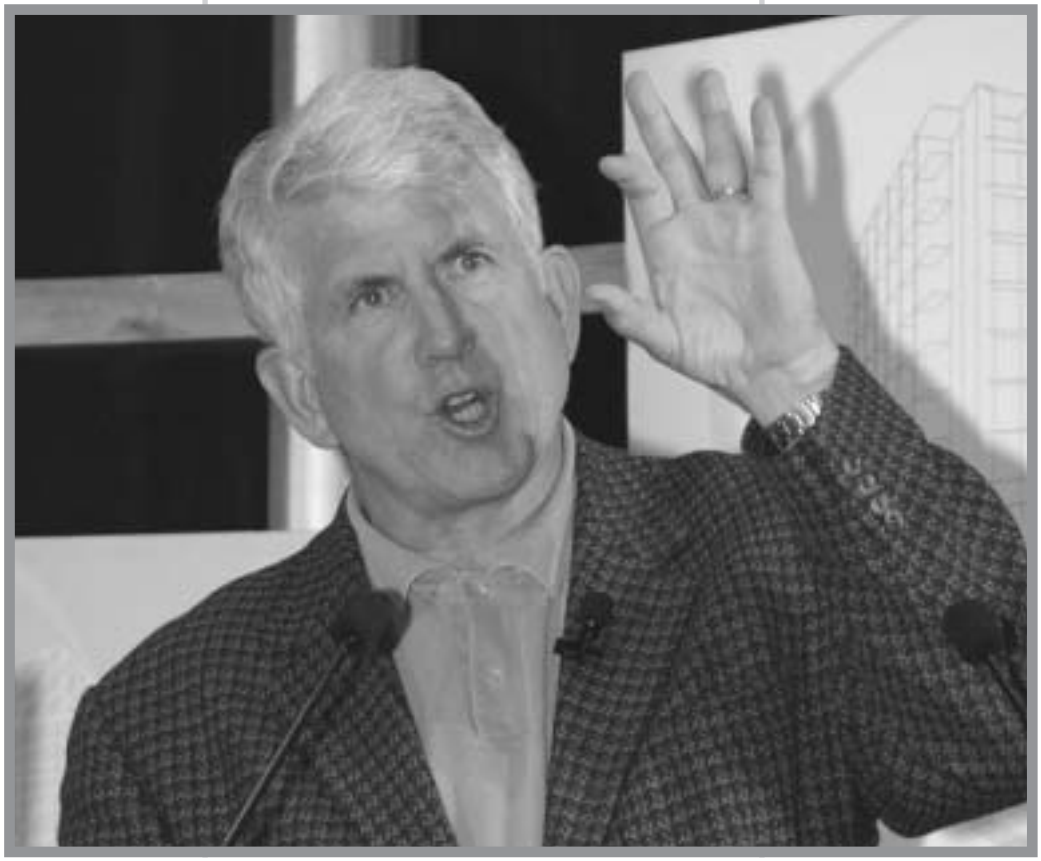

of protocols. You may remember the ISO reference model. Does anyone remember that model with seven layers? That came out of the ARPANET. It was the idea that we were going to take all the stuff that has to be done to do communications and put it in layers and break it up, make a separate committee for each one—that was a mistake—and then design each of the layers independently so that they might operate and evolve independently of each other, which they have. Ethernet was in layers one and two, and in that layering model we were able to make Ethernet simple and robust and cheap. It didn't do everything so people used to say that Ethernet is unreliable and insecure. And it is. It's just that Ethernet relies on the higher-level protocols to take care of all that. Meanwhile, our opponents were trying to build security and other fancy stuff into their local area networking hardware, which caused them to be very expensive and late. That's why Ethernet killed them.

You rely on resources at the edge, and as the edge grows the resources grow and the network scales.

The third good idea was decentralization—another idea stolen from ARPANET and ALOHA. In essence, you want to keep stuff out of the middle, and by keeping stuff out of the middle, you get reliability and scalability. You rely on resources at the edge, and as the edge grows the resources grow and the network scales. For example, on the Ethernet the sharing of the transmission medium has no central control. It relies on randomness; it also relies on control theories so that the randomness leads to a somewhat stable and efficient throughput. Now you can centralize the resources in an Ethernet world. So, for example, Ethernet started as a bus and today it's a hub, but the fact that it can be completely decentralized and have a completely passive and almost empty core is its great strength and gives it much of its robustness. It can be centralized in some cases, but it can also be completely distributed.

The fourth good idea came in that memo on May 22, 1973, in which I chose to use the word "ether," a writer even then. When I came to the point when I had to say, "We're going to send the packets into the coax. We're going to transmit the packets into the coax, which will carry them all over the building," I stopped and I said, "Coax. No, no, no. That's too limiting. We have to call it something that's not so specific, something that's general that gives room for evolution." So I chose the word "ether." We send packets into the ether and they propagate through the ether, which is an omnipresent, passive medium for the propagation of electromagnetic waves. That was a great idea because then Ethernet went from the thick coax that we originally chose to a thin coax, which allowed it to be sold cheaply for personal computers to twisted pair shielded and twisted pair unshielded optical fibers and now all the way back to ALOHA Network through the 802.11 Wi-Fi wireless Ethernet, as I'm fond of calling it. It also allowed the Ethernet to evolve from a bus to a hub, from shared to switched, even to point-to-point to go from 2.94 to 20 to 10 to 1 to 10 to 100 to 1,000 to 10,000 megabits per second, all, I think, just by having chosen that word. In contrast, by the way, I chose the world "collision" to describe when two packets are sent at the wrong time—that was a huge mistake calling it a collision—because most people know what a collision is; it involves broken glass, bent

metal, and a trip to the hospital. The fact that the Ethernet had collisions made it hard to sell during the eighties, but ether was a good word choice.

The fifth winning idea is Ethernet's business model. It's not unique to Ethernet, but has certainly been perfected there. This business model, which I contrast to the vertically integrated IBM business model and to the Cisco, Intel, Microsoft models and to the Open Source model, is the Ethernet model, which goes like this. First, it begins with a *de jure* standard—a standard developed in a legitimate standards body. Contrast this to a *de facto* standard like Windows or a *de ibmo* standard. Have you heard of the *de ibmo* standard, you Latin scholars? Well, I made it up. I'm following in your footsteps, Nick. I just made it up: *de ibmo*. A *de ibmo* standard is neither *de facto* nor *de jure*. It hasn't been shipped yet, but everyone knows it's going to be the standard. And the last *de ibmo* standard by IBM was its Token Ring, which Ethernet famously killed during the late eighties. The Ethernet business model begins with *de jure* standards, which are very painful to make, like making sausage or laws, but they're worth it in the end. The second feature of the business model is that the implementations of this standard are owned by a modern technology corporation, or many of them. This is unlike the Open Source model where you're obligated to give away everything, hoping to make money some other way. Then what follows is the third feature of the Ethernet business model: fierce competition. This competition is on availability, integration, delivery, price, density, form factor, but it's not competition on one thing, which is the fourth feature of the business model: interoperability. In the Ethernet market, competitors are not allowed by customers to compete on incompatibility. I'm reminded of the Ethernet transceivers we first made in 1981. AT&T wanted us to build 14-megabit-per-second Ethernet transceivers because I would give them room to be up to 40 percent faster than the competing products. Duh … so we built a transceiver that could run at 10 and 14; it never ran at 14, of course; it always ran at 10 because AT&T wanted its computers to talk to other computers. And the fifth feature of the Ethernet business model is rapid evolution in the *de jure* standards body, based on market engagement and technological innovation. You have all these versions of Ethernet that have come out and are still coming out to this day, which leads to the sixth feature of the business model: that evolution is constrained by an emphasis on leveraging the installed base. When we went from 10 to 100 megabits per second, you could buy a 10-100 card that could talk to a 10- or a 100-megabit switch, and the install base was thereby preserved. This is where Metcalfe's law kicks in. If you preserve that installed base, then you still have the value of the network even with upgraded capabilities. So the five ideas that were winning with Ethernet, including the six features of the business model, should be handy in the next thirty years, especially here on earth.

I just told a joke and only George laughed! I'll have to try again.

So now on to some broader … are you still there? I just told a joke and only George laughed! I'll have to try again.

Ethernet was invented at Xerox Park; I think I've mentioned that. It's an example of an invention, a nice one that came out of a laboratory that was supported by a monopoly. I've come to the conclusion over the years that only monopolies can afford basic research.

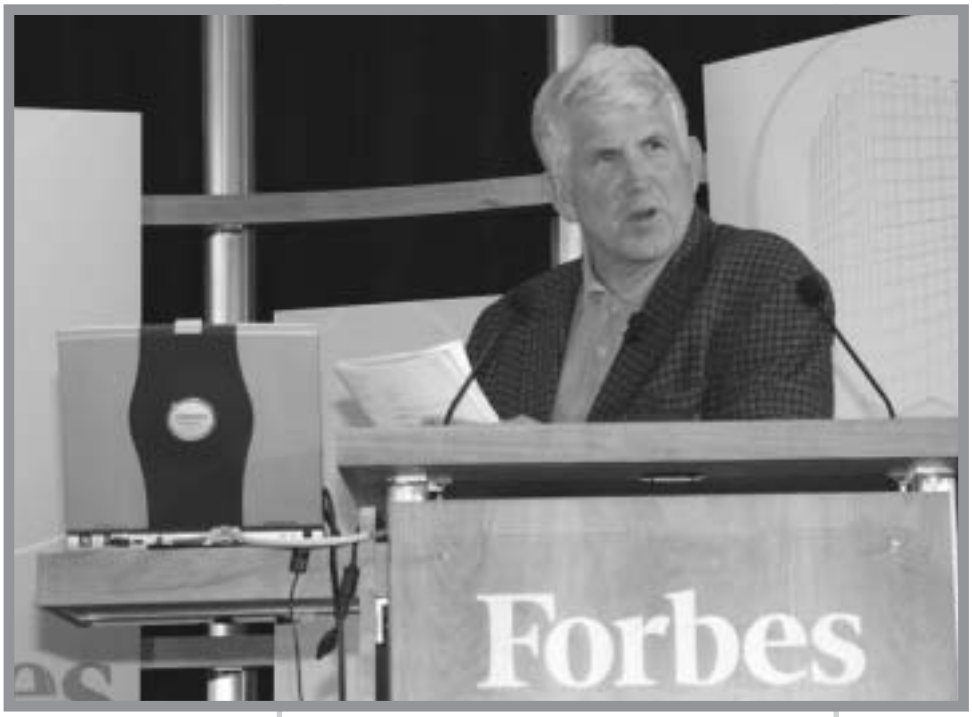

AT&T, for example, up until 1984 was this massive monopoly that was able to fund Bell Labs, and, of course, in return for that monopoly out of Bell Labs came the transistor in 1947, the Princess telephone in 1957, and UNIX in 1967 or 1968. IBM was a monopoly for a long time, up until around 1984 when they had the Watson Research Centers, which also gave us many wonders including punch cards, Winchester disks, relational databases, etc. But that leads to the question: is it worth tolerating monopolies so that they can have research centers to do basic research so we can get all these cool things like Ethernet, like UNIX? And I, at my ripe old age, have an answer to that question. No!

No, it is not worth it. Those monopolies were not worth all of that technological innovation. Especially monopolies, by the way, created by the government. So I claim that what we want to do is kill these monopolies. I agree with you that we should abolish the FCC, but we should only do so the day after it succeeds in killing Verizon. Is Verizon one of your sponsors?

Anyway, I've become a big fan of FOCACA, as I mentioned in my introduction. FOCACA stands for Freedom of Choice Among Competing Alternatives. When FOCACA reigns, everything goes well. I'm bringing this up because it's the only thing I disagree with George Gilder about, and I'm worried sick. I really worry when I disagree with George. One of the purposes of government is to protect people from violence—national security, law enforcement. Well, there's such a thing as economic violence, and monopolies when they engage in anti-competitive practices are engaged in economic violence, and it is the job of government to stop them.

George Gilder:

Economic violence is [MIT professor and prominent political dissident] Noam Chomsky's favorite concept (much laughter).

Bob Metcalfe:

Well, there goes that one! Now I know why I was worried. Is my time up yet?

So really, where should basic research be conducted? Basic research should be conducted at research universities, and they should be funded by the ultimate monopoly—government. Government is the monopoly with the M1A1 tanks. Universities are the best place to produce research because they produce people, and those people graduate and enter the world with the technologies that they've developed. That is the most effective technology transfer mechanism. It does

not work generally speaking to throw technologies over the wall; it really works to carry them with you into commerce. Now if you accept that, that puts the onus on us to do a much better job of managing our universities, which are very badly managed. They have this thing they do at universities. If you see something completely ridiculous going on you say, "That's completely ridiculous." And they say, "Bob, you're a business guy. You just don't understand how the university works." I've had enough of that so I'm making trouble at my favorite university to keep it at the top of its game as a place where basic research is done and good people are graduated. Now one of my favorite expressions is that invention is a flower and innovation is a weed because the invention at Xerox Park was conducted in the hot house where we were allowed to live in the remote future and build networks for PCs, even when there weren't any PCs in the real world. But once I left Xerox Park, I found that there was no hot house for orchids out there. The world of innovation is a tough and ugly place, and the reason is that the status quo is very resilient, and the status quo will lie, cheat, steal, and kill to keep its position. When you're an innovator and you've got your invention and you go out into the world, it does turn ugly pretty quickly.

The world of innovation is a tough and ugly place, and the reason is that the status quo is very resilient, and the status quo will lie, cheat, steal, and kill to keep its position.

I'm proud to be a citizen of the U.S. portion of the Telecosm. I'm grateful for this opportunity to have tried out some of my ideas, and I'm enthusiastic about the future but I have just one problem that I'm left dealing with. Based on the thirty years of the success of Ethernet, I am now the status quo and I don't really know how to behave. I'm trying not to be one of the monopolistic deadheads who run the status quo. I am trying to behave myself, but I really don't know how. This is a whole new experience for me.

Steven Sprague, President and
Chief Executive Officer of Wave Systems

The All Optical Network

tuesday

august 26

3:00 pm

George Gilder:

I have high hopes that David Huber's Broadwing network can create a monopoly of innovation.

David Huber:*

Thank you, George. I first became aware of George when I was in the process of starting my first company, Ciena, which was the first company to bring wavelength division multiplexing (WDM) to the commercial telecom market, and George had been writing about the possibilities of the wavelength, which would bring ubiquitous bandwidth. I think among all the business writers who write about technology, George is focused in the right direction and understands the implications. What I think we've all seen here recently is what we're not quite so sure about is the time frame for all of this, but I would say we're very certain about the end result. I'm very happy to be here, particularly because it's George's conference, and I have a great respect for George and his view of how technology will affect business and everyday life.

* Chairman and CEO, Corvis Corporation

We started Corvis in 1997 and brought our product to market in 2000. Our timing was the best and the worst. We hit the peak of the equipment valuations, but we also hit right at the point in time where carriers had put the brakes on spending for new infrastructure. What we didn't know in the year 2000 was how severe the spending decline would be. As we took a look at the company eighteen months ago, we realized that because there weren't a lot of new builds to win we needed to look at some other areas. And, of course, we're very big believers in optical networks to provide bandwidth for the future. In fact, I was very interested in some of George's recent writings in which he points out that Korea consumes much more bandwidth per individual than we do here in the United States. So I think we have some issues there, and there are some real benefits to inexpensive bandwidth. Another thing we've been happy about recently is that the U.S. government has made a very strong statement that they want to build a network to help them with the national security needs of this country, and they have made it extremely clear that they are focused on an optical network. They're the world's largest consumer of telecom services; they're very experienced with the SDH and SONET networks, but that is not what they're looking for today. Today they're looking for an all-optical network. Just as the military is making decisions that will affect the security of this country as well as the cost of military data networks and their ability to meet any type of security need or threat that occurs in the future, telecoms are facing similar issues today, and the decisions they make will affect their financial viability in the future.

I'd now like to give a brief review of these optical networks and compare them to the traditional SDH and SONET networks. In these long-distance networks, there are two key parts: the transport piece and the switching piece. In the traditional networks, the switching is all done with electronics. In traditional SONET or SDH networks regeneration takes place in the transport element of the network. Every 300 to 400 miles complete conversion to the electronic domain is required. You may ask, "What's wrong with that?" Well, there's nothing wrong with that other than the fact that it's extremely expensive and it decreases the reliability in telecommunications networks. But, the bigger issue is the switching fabric. Typically, electronic add/drop multiplexers form the core switch fabric in the long-distance networks. The problem is that the data gets switched over and over again. It gets switched about every few hundred miles, which leads to very high costs not only in capital but more importantly in the personnel it takes to run these networks. In the early part of my career, I was at Rockwell International in Dallas when the SONET standards were being developed. It's actually a great standard for voice but it didn't anticipate the growth we'd have in the data networks. It's a technology that was designed for voice but doesn't serve us particularly well in the data world. You've heard about stacked rings in the SONET world where one wavelength is not enough so you continue to add a bunch of little networks on top of the original network

I don't think many people or much of the business world understands what optical networks mean and how they can be of benefit.

that was designed with one wavelength in mind.

I'm not sure how great an evangelist I've been because I don't think many of the business world understand what optical networks mean and how they can be of benefit. We can compare SDH or SONET networks, where the switching takes place every 300 miles, to an airline. Suppose you need to make a trip from California to New York and you're offered a ticket on an airline that sends you through a hub every 300 miles. At each hub you'd have to change planes and continue your journey. I don't think we'd have any volunteers because it would take you several days to make the trip, and yet that's exactly how the SDH or SONET networks work. They go through the add/drop multiplexer every 300-400 miles, and all the traffic is again switched and rerouted. Just as you wouldn't take a trip like that on an airline, it's no way to run a network. We're moving data in the large volumes we are today and in larger volumes, I hope, in the future, so Corvis can maintain its competitive lead in the industry and our standard of living in this country. In the networks that Corvis builds, specifically in the Broadwing network that we purchased two months ago, if you need to go from Los Angeles to New York or New York to LA, those are always nonstop flights. We don't go through electronic switching every 300 to 400 miles, and we don't convert from optical into electronic domain very often because of the very high cost and lower reliability we receive in the networks. What we end up with is an all-optical network with all nonstop flights.

I can't resist the temptation to point out the dirty tricks that our competitors play. One of our major competitors claims, "So and so, maker of all-optical switches," but if you get down into their "optical switch," you'll find it only contains electronics. When Corvis and Broadwing talk about optical switches, we talk about switches where photons go in and photons are routed. You can think of an optical switch as a switch with a mirror that is color selective and makes switches based on the color of each wavelength of light. And, no, electronic switches are not optical switches. They may be connected to optical fiber, but that does not make them optical switches. In that regard—and this will be the last remark of this type—throughout the telecom boom of three or four years past there were about fifty companies that tried to make optical switches for the core network. I might add that some of the largest names in telecommunications tried to make these devices and failed—in most cases from a technical point of view, but in many cases for business reasons—to bring these to market. Corvis built these machines, brought them to market, and I'm very pleased to say—I'm very surprised by this actually—that after three years of operations in a number of commercial telecom networks, we've had no failures with our optical switches. That is very surprising for a new technology, considering our competitors weren't even able to make them work.

I can't resist the temptation to point out the dirty tricks that our competitors play.

How do you make an optical network? You need ultra-long-haul transmission. Why do you need that? Without it your plane runs out of fuel and you have to stop for a landing in that cumbersome process we call electronic regeneration. Erbium-doped

amplifiers (EDFAs) were a great step to get us going with wavelength division multiplexing, but they didn't offer the long distance. You need the Raman amplifiers. Then you need the optical switch so you can go through those hubs without actually landing the plane. You just go through it on a routing basis without actually unpacking the plane and changing the packets around. And then, of course, you need changes in the way telecom traffic is groomed and provisioned. Packets should be groomed (sorted) at the beginning and at the end. Don't groom on an intermediate basis.

Why are optical switches important? Because optical switches cost forty to a hundred times less on cost per gigabit than electronic switching. We need to make the switch in telecom networks from the networks about voice to today's business, which is about data. A lot of the turmoil that you see in the telecom industry is because of this switch to data. The economics are very different.

With EDFAs, you arrive at the hub and the amplifier changes design so that the solo noise of the signal is still high enough that we can amplify it and move on to the next hub, which would typically be 80 to 100 kilometers down the way. Then it needs to be amplified again. With a Raman amplifier, you still have the peak of your signal at the hub, but unlike the EDFAs, the signal begins to be amplified before it reaches the next hub. How do you start to amplify the signal before you reach the optical processing unit? In Raman amplifiers you can actually project an amplifier out into the standard fiber and use a standard telecom's fiber as an amplifier. What the Raman technology does is it effectively moves the amplifier spacing from that 100 kilometers down to 80 kilometers. In doing that, you can change the requirement for regeneration, as in the traditional SONET networks, from about 240 miles up to 2,000 miles simply with this small change in effective amplifier spacing. Corvis is very proud of its achievements with the optical switch. All the optical switch does is route traffic in a city based on the color of the lights. A freeway system that has multiple freeways coming into one spot is called a mixmaster. When traffic is coming into one optical switch from all different geographical points, it re-switches the traffic to its final destination based on the color.

Optical switches cost forty to a hundred times less on a cost per gigabit than electronic switching.

Why do you make an all-optical network? Number one, lower operating costs: fewer people and faster provisioning. Take a look at the financials that have come out of Broadwing and will continue to come out of Broadwing. You'll find that the revenue per employee of this very small network, which doesn't enjoy the economies of scale of some of its largest competitors, is already twice that of some of the biggest carriers in the world. You'll see that continue to be demonstrated. Two years ago Broadwing won a very intense competition in putting fifty cities up in the "carriers' carrier" deal. No other carrier could match them on the speed at which they could deploy these additional cities. The reason for doing an optical network is lower capital costs. In a network carrying a lot of traffic you can't afford large electronic switches everywhere. In three years of

Chief Executive Officer of Essex Corporation Terry Turpin
with Corvis Chairman and Chief Executive Officer David Huber

carrying commercial traffic, none of our switches has had a single failure. Network reliability increases because of simplicity.

The Broadwing network is the world's only all-optical network. Three or four years back, you may recall that Lucent and Corning were extolling the virtues of new lower-dispersion fibers because they could produce bandwidth less expensively. Fast-forward to the year 2003; which carriers has an optical network lit with this fiber? None of the large carriers have even lit the lower-dispersion fibers and of the smaller carriers that have lit it, only one of them has the all-optical network that allows the bandwidth to be produced at the lowest possible cost. If we are to "move up" in bandwidth consumption the way that some Third World countries have we will need a number of optical networks like Broadwing is running today. There are lots of applications for this, and we hope that we move on to enjoy them.

Corvis is a company with the same vision that it was founded with, the vision to develop all-optical networks.

Ashby Foote and Jeff Stambovsky strike a pose.

If David and Mike knew about the open bar they would have signed up for Telecosms I - VII.

Jim and Jennifer Mullens catch up with Ann and Terry Turpin.

Bret, Nick, and Sandy set up camp in the Gilder employees' section of BullWhackers Pub.

Having communicated for years on the Gilder subscriber message board, Roger had always imagined Gene would be much taller.

Richard threatens to put Mary in a headlock if she makes him sit at kids table again.

Carver Mead*

Back to the Future

the analog age

tuesday

august 26

4:30 pm

George Gilder:

Foveon's saga is really the life story of Carver Mead, which I first told in my book on the semiconductor industry, Microcosm. (See also "The Spectator Interview: Carver Mead," *TAS*, September/October 2001.) A pivotal point came in 1986, when the Valley was roaring back from its last great cataclysmic slump, with revenues dropping some 45 percent in a year. In a Caltech classroom in Pasadena, the eminent Gordon E. and Betty Moore Professor of Engineering and Applied Science, like many in his trade, seemed to be flaunting his august connections to the technologic eruption underway up north. Projecting the design of a massively parallel processor on the screen, he proposed it as a model for a revolution in computing and said: "Now I've been up in Silicon Valley, talking to the guy who made this thing and …"

Why is this class laughing? Don't they believe in Mead, the industry's first and most profound prophet of Very Large Scale Integration—VLSI—microchips? An intimate of many of the founders of the digital age, from inventor of the integrated circuit Bob Noyce to microprocessor architect Federico Faggin? Indeed, Mead had taught them much of what they knew about the design of digital devices. He had performed the crucial researches from which Moore's law itself derived, ordaining the doubling of digital computer performance every eighteen months. But the design he was showing on the screen to such friendly hilarity was not a digital device at all. It was analog, not a bit or a byte in sight. It was a schematic of the human brain.

Whether it was God or Gordon Moore, whom Mead had been consulting up there in the Valley, Mead's citation of the brain was not unusual in computer science. What was radical was that rather than treating the digital computer as a possible model for an extended or ultimately superior brain, he was offering the analog brain as a model of an incomparably more powerful computer. After twenty years as the industry's most authoritative proponent of the power of digital electronics, he was reversing direction and declaring the onset of a new era, the analog age.

Mead's analog technologies will change the world, but the world will have to change, too, to accommodate these new capabilities.

*** Chairman, Foveon • Gordon Moore Professor of Comuter Science, Caltech**

Carver Mead:

"To understand reality, you have to understand how things work. If you do that, you can start to do engineering with it, build things. And if you can't, whatever you're doing probably isn't good science. To me, engineering and science aren't separate endeavors."

"Research is a matter of love. It's not a left-brain thing. Once you figure out something, then you construct an elaborate rationale—the talks you eventually give that make it all sound so simple."

"The problem with 'demand pull' is that by the time you have a real product, the market will have moved on. You're doomed to playing catch-up. I prefer 'technology push'—find an interesting new technology and try to come up with uses for it. 'A solution looking for a problem" is supposed to be a terrible epithet, but in my experience it works.'

"Impinj, a company started by a former student of mine at Caltech, Chris Diorio, started out with something completely unrelated—neurally inspired computing—and came up with a very precise and low-power way to put a charge on a floating-gate transistor, which is the basis for flash memory. It was a classic 'solution looking for a problem,' which is turning out to be RFID, the little [radio frequency] identity tags to put on things. They're the ultimate lower-power device—picowatts, whatever you can get out of a little antenna. So instead of just having a 'dumb' tag that can tell you its name and nothing more, you get a smart one that updates itself as it goes. You get a package or a product that can tell you its whole history, right there."

Carver Mead and his student, the late Mischa Mahowald: "In digital systems, data and computational operations must be converted into binary code, a process that requires about 10,000 digital voltage changes per operation. Analog devices carry out the same operation in one step and so decrease the power consumption of silicon circuits by a factor of about 10,000."

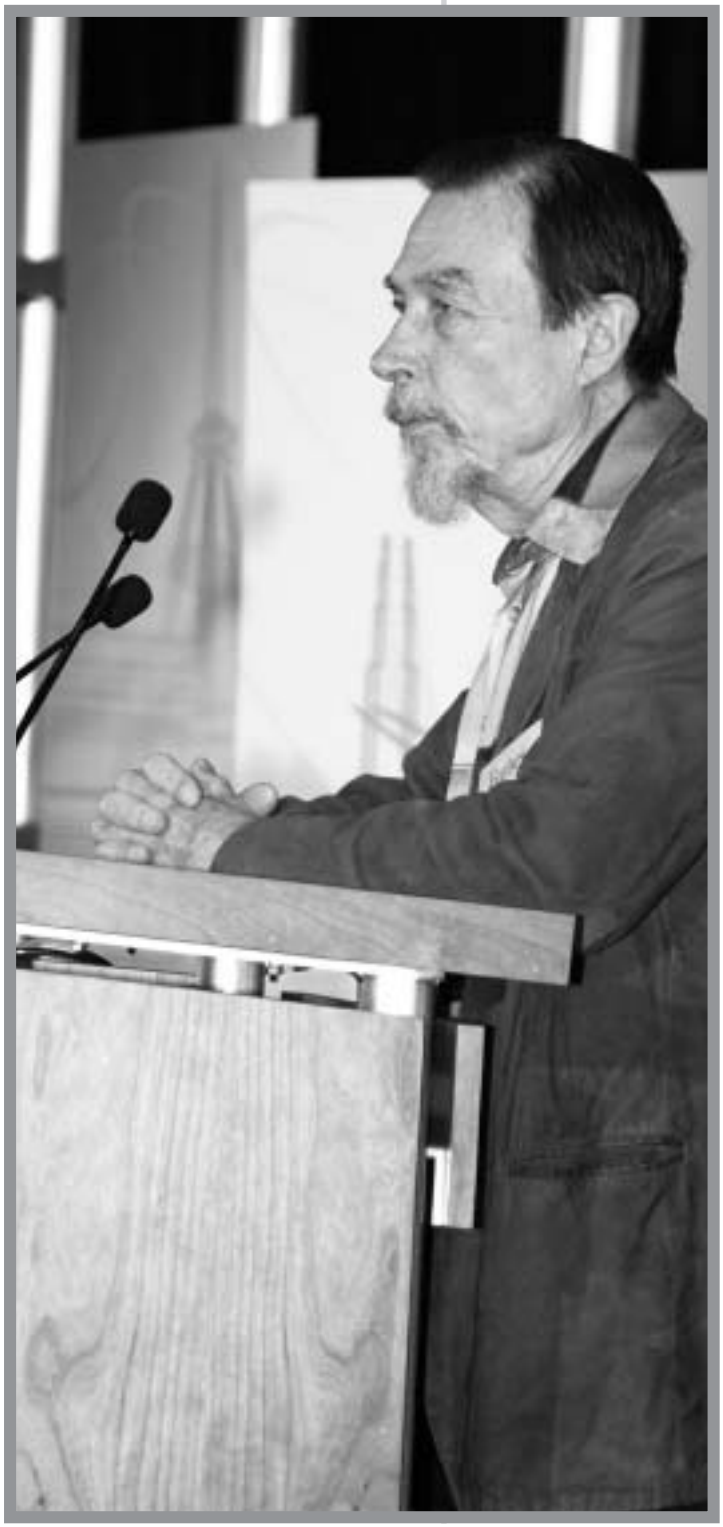

Telecosm 2003 Acronym Decoder

ADC	analog-to-digital converter
ASDL	asymmetric digital subscriber line
ASIC	application-specific integrated circuit
ASSPs	application-specific standard products
BREW	binary run-time environment for wireless (Qualcomm™)
BWA	broadband wireless access
CDMA	code division multiple access
CDMA2000	third-generation CDMA
CPU	central processing unit
CDSL	consumer digital subscriber line
CLEC	competitive local exchange carrier
DHWG	Digital Home Working Group
DSL	digital subscriber line
DSP	digital signal processor
EDFA	erbium-doped fiber amplifier
EvDO	evolution-data only
FIM	fault interceptor module
FPGA	field programmable gate array
GPS	global positioning system
GPRS	general packet radio service
GSM	groupe special mobile
IP	Internet protocol
ITU	International Telecommunications Union
JFET	junction field effect transistor
LAN	local area network
MAN	metro area network
MIDI	musical instrument digital interface
PCA	Intel term for personal client architecture
PLD	programmable logic device
RAID	redundant array of independent disks
RBOC	regional Bell operating company
RFID	radio frequency identification
S-CDMA	synchronous code division multiple access
SERDES	serialization/deserialization
SDH	synchronous digital hierarchy (networks)
SoC	system-on-chip
SONET	synchronous optical network
SMR	specialized mobile radio
SMS	short message service
TDMA	time division multiple access
USB	universal serial bus
WAN	wide area network
W-CDMA	wideband code division multiple access
WDM	wavelength division multiplexing
Wi-Fi	wireless fidelity (used when referring to 802.11a, 802.11b, dual-band, etc.)
WiMAX	802.16a wireless metropolitan-area standard
WPA	Wi-Fi protected access

Raise your hand if you thought Telecosm 2003 was the best Gilder/Forbes Telecosm. Ever.

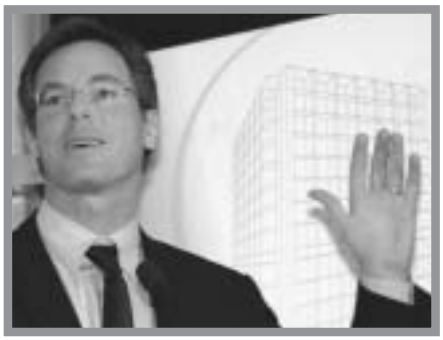


ALTERA.


cloud


corvis


endwave


GILDER
TECHNOLOGY
REPORT


KENT


QUALCOMM


SAMSUNG


verizon


WIDERAY