RealComputing NPR Radio Transcript

Week of September 9–16, 1998

John C. DVORAK : Hi, welcome to Real Computing. I'm your host John C. Dvorak. Every week we take a look at the world of personal computing, high technology, the Worldwide Web and everything in-between, with this week being no exception. We've got Carolyn Rose, CEO of U.S. Web Learning here to talk about some of the new programs that that company is developming to help people learn more and more about the Internet and how it works, mostly to become professionals within the business. Ami Amir from RadVision out of Tel Aviv is here. He came all the way from Israel to discuss h.323, which is the video conferencing standard that you find in most software nowadays, and how it came about. Also, some of the new stuff that the company is up to. It's a very interesting story because h.323 came in out of the blue and kind of solved a lot of problems that we were all having trying to use video phones, for example. And finally, Dr. Neil GUNTHER, is an author and professor and expert on PC performance, is going to tell us what's wrong with the industry, right after the news.

. . .

Maureen MCGINLEY : And now, here's your host, John C. Dvorak.

DVORAK : We have with us Neil GUNTHER, who is an author, professor, and one of the world's experts in PC performance. In fact, one of his books The Practical Performance Analyst is a famous text from McGraw-Hill that people will use to analyze the various performance elements of the world of the PC. Neil, welcome to Real Computing.

Neil GUNTHER : Thank you.

DVORAK : Now, you do a lot of presentations about some of things you dislike about the PC world, and I guess performance being an issue.

GUNTHER : Yes, performance is my major focus.

DVORAK : What is the state-of-the-art in terms of... you know, what's a good machine out there, in terms of the high performance? Is there anything out there that's worthwhile?

GUNTHER : Oh, I think almost every machine is worthwhile, including the new iMac.

DVORAK : You like that machine?

GUNTHER : I haven't used it, but, you know, I'm sure they've done a good job on it. And I use a Mac a home. I have a Power PC, I'm a Mac person.

DVORAK : So you're a Mac bigot.

GUNTHER : Because I worked at Xerox PARC for eight years, I was imbued with that kind of technology that Apple helped to get into the world. And so I enjoy using it. But my point is that it's not so much is there a machine that has the kind of performance that meets your needs, my take on a lot of this is that there's more performance out there than you can even use. And the only reason you have to upgrade is because of things like continued software bloat. You know, and the way to look at this is that when we do our work everyday–I do a lot of work on computers just from an office point of view–I'm not really doing anything very different from what I was doing in 1980. But in 1980, I could use a relatively small CPU, running just at 3 or 5 megahertz, and I only had a half-a-meg of RAM. I could still do word processing, spreadsheets, and other things that I need for doing classes and textbook preparation–something that I do. Why do I need to have now 40 megabytes in my Macintosh, you know, running at 66 or 100 megahertz? And now I still don't have enough memory to run Netscape, Word, and Excel all together; it struggles sometimes. And the reason for that is because I can't switch Virtual Memory on in the... as Virtual Memory. Most applications value to turn it off because the performance is better without Virtual Memory. If that's the case, why did Apple bother developing it?

So if you stop for a moment and just look at what's going on, particularly from a performance perspective, you have to ask yourself why are

we on this upgrade treadmill. What is the performance of it? Because we've got more performance than we really need.

DVORAK : Well, you know, that's a funny thing, because I just finished a column on this, as a matter of fact, which is that a microprocessor report is showing in their latest edition a tendency for the demand curves to cross over from available CPU power and necessary CPU power. In other words, how much power do you need...? The treadmill problem. And apparently, we've gotten to a point where because of hard disk capacities and some of these other things have spurted ahead of Moore's Law, that we don't need anymore... I mean, people are sensing that they don't need to upgrade. Because as hard as Microsoft tries, they've been having trouble making the bloat worse than it is.

GUNTHER : Right. And your assistant was telling me before than Intel is now engaged in trying to get people to buy, I think, it's the MMX technology via the Web browser. When people use Intel advertising in their websites, the have to have it done specifically, it has to be highly graphical...

DVORAK : Well, they're trying. Nobody's doing it. I mean, why would I want to... Because I know there's such a small audience. If you want people coming to your website and you want to disgust them with a slow-performing website, yeah, sure, you can do that. But take a look at the company... the only company making money on the Web is Yahoo. They've got nothing fancy on their website.

GUNTHER : If I can draw an analogy with Moore's Law—Moore's Law is an interesting one, let's take that for a moment. Let's think of that in the automobile industry. Let's wind the clock back to 1980. So we've got roughly a 20-year period there. Now, Moore's Law says that CPU and memory, speed, and design is increasing roughly by a factor of two, they're doubling it roughly every 18 months to two years. If you go back and take a car in the 1980's and wind that forward, using the same scale factors that apply to a CPU and memory chips and so on that's going on in the CPU industry, if you do the calculation, it would come out something like this. The current automobile you have would have something like 1,000 horsepower, and it would have a 150-gallon gas tank, and it would weigh something like 10 tons. And that's because

they've added all this stuff on, all these additional features into the car that now drives like an ARMY TANK. And it's still questionable whether you need it, because all you really want to do is to get to work on time, or you might want to do your shopping, or you want to go on vacation, do some traveling. You do not need 1,000 horsepower to do that. If that's unacceptable in the automobile industry, why is it acceptable in the PC industry, to go on upgrading, getting more CPU power, when you can't even use it?

DVORAK : Hey, that's the question a lot of people are asking. But it looks like it may answer itself. You also have had scathing reviews of Windows NT.

GUNTHER : You mean what I've written?

DVORAK : Yeah, have you?

GUNTHER : I wouldn't classify them necessarily as scathing, but I've taken a...

DVORAK : Skeptical perspective?

GUNTHER : Yes.

DVORAK : Why?

GUNTHER : Well, because I'm trained as a skeptic, as a scientist.

DVORAK : So what's wrong with NT?

GUNTHER : I don't think there's anything wrong with NT. My reaction is more to the way Microsoft has been presenting Windows NT, and software in general, as, you know, God's gift to the computing industry. And they've just simply gone completely overboard. And I've written a couple of articles in response to that, because I've tried to counterbalance some their hype for those people who don't understand the kind of tactics that they're engaged in at the marketing level.

DVORAK : Well, let's go over a few of those.

GUNTHER : Okay.

DVORAK : For example, I think one of the things you commented on was scalability.

GUNTHER : Yes. Well, scalability is one of my little pet things, and, you know, it's chapter 6 in the book, or something like that. Basically, one of the things that Microsoft has been saying is that, you know, we can only produce machines... this is really now more for commercial use than the PC industry. But when you have commercial machines, you need to have more than one CPU to run, for example, something like a big data base, like Oracle, for example, or Microsoft SQL server. So ultimately, what you need in order to be able to handle multiple users, large data base, in other words, large scale systems, you have to have more than one CPU in the system. Now, this has been done in the UNIX arena now for 15 to 20 years. And in the UNIX world, it's well understood how to do this. Now, they've basically refined it and these machines run very well in scale, very well in general. And Microsoft has had the audacity to attack that on the following basis. They say, Well, we can only do four CPU's now.' This is roughly where it stood about a year ago. We can put four CPU's in a box. And our scaleability is linear–as you add more CPU's, you get equivalent bang for the buck. So if you put four CPU's, you can get four times the amount you have with one CPU.' In general, that low CPU number, that's approximately true. But as you go up into more CPU's, you tend to get a drift bandwidth from linear, in other words, the curve falls away. This is classically known. This is basically an example of something called Amdahl's Law. You talked about Moore's Law before, this is another one, Amdahl's Law. So this is well-known in the industry. And the question is: how efficient can I make that curve as it falls away from linear? Because I'm going to have to suffer that over-hit, to have these separate CPU's orchestrated and running together. But Microsoft says, Well, we're not there yet. We can't do 64 CPU's, for example, like Sun Microsystems can running UNIX. But we can do four. And if you look at our... curves, they're fairly linear, so therefore, by induction, the curve would continue to rise in a linear way. Therefore, our scaleability is better than most UNIX platforms.' Well, of course, this is nonsense. And, in fact, I would go even further and say the reason they can't scale beyond four–and they're currently, about roughly, six to eight CPU's–is because–this is the technological regression point that I mentioned

to you earlier–that in the PC industry, what they try to do is put together systems that are made from the cheapest components–that's the commodity marketplace, it's a cutthroat business, it's not to do with technology. And the consequence of that, the technology actually is not the best you can get, it's whatever the market can bear in terms of cost. And that tends to be something regressive from what's the best available, because the best available is not the cheapest. So in order to put these PC chips together to make, for example, four, six, or eight multi-processor in one box, you have to have a bus that connects them so that they can all talk to one another. But the bus efficiency then determines that scaleability, and a very inefficient bus means that the scaleability curve would fall away from linear very quickly. And, in fact, that's what does happen, because the P6 bus for the Intel chip is relatively inefficient. Now, this is not because these people are silly, it's just because they don't have the experience. They're coming from the desktop marketplace and trying to move up into the enterprise. In fact, UNIX has gone down that same road, except they're about 20 years ahead, roughly. And this is all just sort of a chicken and egg thing, if you look at it, because what UNIX did to the mainframe, the PC world and the WinTel game will now try to do to UNIX. So there's nothing particularly new here, but the WinTel platform is relatively immature. And the reason they're immature, partly, is because they're using PC technology, that is to say, the cheapest commodity technology, which does not have the best performance, does not have the best scaleability, and it's going to take a long time to get there. Microsoft, on the other hand, says that this is blazing the obvious, that they already have it in spades. I'm simplifying the argument because I don't want to get into all the details, but if you look at that argument in detail, you find it's really quite bogus. And I object to that because it's been pushed forward by people who work for Microsoft, who previously would never have said such things. So they have apparently struck a bargain with Redmond, Washington, and now they're saying things that, really, they would not have said before and that are technically quite incorrect and very misleading.

DVORAK : Well, I notice that a lot of people are jumping on the NT bandwagon, but there's been a lot of disappointment recently because of this long delay of Windows NT 5.

GUNTHER : Well, that's an interesting point. NT 5, if you want a little perspective on that, Windows NT 5 will roughly be, particularly from the point of view of measuring performance and so on of the system, will roughly bring it up to where UNIX was, let's say, about five to ten years ago.

DVORAK : So if you were consulting with somebody, you would still be recommending UNIX for simple servers and things like that?

GUNTHER : It depends what they're trying to do. I wouldn't necessarily make a flat-out recommendation like that, but what one would have to look at is if you want high capacity, high performance servers to run large data bases and so on and so forth, then more than likely you're going to end up looking at UNIX servers because they have the maturity, they have the scalability. It's just a fact of life.

DVORAK : How did you get involved with the PC performance issues to begin with?

GUNTHER : Well, I'm really trained as a theoretical physicist actually.

DVORAK : Oh, one of those guys.

GUNTHER : That's right. And, you know, in order to try and make a living, you find that you have to go and do something different ultimately. So I've spent eight years at Xerox PARC. And that was really my computer education. I saw the Macintosh long before the public saw it. And Xerox was using that internally in amazing ways. For example, I worked in the chip fab line at Xerox PARC, and the whole process line was run by a series of Altos, which you can think of as the Macintosh from 1974, not 84, built by Xerox. Or connected by the Internet to file servers. And each process step that was done on a wafer was communicated to the next engineer on the process line through a piece of software that ran on the top of their e-mail system–very, very cleaver stuff. Ultimately, I got involved in building the next generation machine called the Dragon. And as a consequence of looking at that, I ran into some questions I asked myself about, you know, how would you load down on such a system? That was going to be a multiprocessor workstation for personal use, very different from what we built before. And as a consequence of asking question about how would you most efficiently use that system,

I just basically kind of fell into this whole area of performance analysis, which I didn't even know existed.

Interesting thing about that is that performance analysis is actually very similar to physical science. The reason is because when you're measuring a system or trying to analyze a system, you have to form a hypothesis about its behavior, and then you have to test that hypothesis against your measurements. Now, that's a physical science process like a physicist or a chemist would do, but computer scientists never learned that. So, in fact, performance analysis is kind of this funny area, this weird area, that's sort of related to computer science, but somewhere between real physical science and computer science, and it requires having a lot of training in physical science-type activities, even though you're looking at an un-physical engineered system.

DVORAK : What happened to the Dragon?

GUNTHER : Interesting question. The Dragon, actually, as you know, Xerox always fumbled getting product into the real world, and same here. Around about 1988, I think, if I remember it correctly, or so, Xerox entered an agreement with Sun Microsystems, interesting enough, to take their technology and put it into a product which Sun developed, and that became the Spark Center 2000 Multiprocessor Server that Sun started building in around about 1990.

DVORAK : Now what about the Star technology? Didn't that derive from the Alto 2?

GUNTHER : You can think of it as being a derivative. And that's a very interesting story from my point of view, because, in fact, the Star was essentially the Xerox version of the Macintosh, roughly speaking. However, it came out about the same time as the Lisa, prior to the Macintosh.

DVORAK : Right, it came out actually just very shortly before, maybe six to nine months. And it costs three times as much, and the Lisa was expensive enough.

GUNTHER : And the Lisa and the Star workstation were roughly about $10,000 apiece, which is way above most people, in terms of the home

user, it was way above their price point. So there was a technology price point problem, and the failure that both Apple and Xerox had is that they thought that when people saw the Window-GUI interfaces we now expect on every PC, when people saw that, they would just naturally gravitate to it. In fact, they were totally wrong. And a very interesting example is I came to the computer fair in 1983 or 4, whenever the Lisa was delivered.

DVORAK : 83.

GUNTHER : 83. And I was standing in a crowd, and Apple had a TV monitor up on rack, and they had the Lisa down on the desk. And all the guy was doing, you can't believe this now, but all he was doing was moving the mouse around on the table and clicking on icons, because there were really no applications that were very interesting at that point anyway. And you could see all the desktop of the Lisa on the television screen. And in the crowd–they were 15 deep–I heard somebody behind me say to their friend, You know, that's not real. What they did was they actually recorded that on video tape, and that's a video tape you're looking at. It's not a real computer.' And I thought that's amazing. Even if that guy's half joking, I think he actually has expressed the public perception that you cannot do real computing with a cartoon interface of the type that was seen on the Lisa and the Star workstation. And that's a psychological problem, that's not an engineering problem. And both Xerox and Apple hit the skids at that point, and neither of those machines sold. And Steve Jobs had to do a lot of work in order to get the Mac into the market, and had some success, but rather limited success, you would have to say, overall.

DVORAK : Yeah, that's true. But it's funny you mentioned cartoon interface, because Microsoft's Bob, which came out, you know, a number of years later, was a very aggressive attempt at doing something different, which, you know, Microsoft's always been criticized for not doing. But every time they do it they fail, and that makes it worse, you know, because then they're more reluctant. But I, when I first saw that, my son at the time loved playing with it, and he did things with it that I couldn't believe were even in this product. I mean, it was actually a deep and interesting product. But it was obviously, like the perfect interface for a kid. And Microsoft didn't get it, because, you know... I

don't know whether it's because they're just kids up there themselves and they don't understand it, or didn't have any children because everybody was in their 20's and, you know, working 16 hours a day, and they didn't get that aspect of life. And so the thing was a horrible failure, because an adult was not going to play with a cartoon interface, which is exactly what it was, and as interesting as it was, when a kid got a hold of it. And they couldn't remarket it any other way and they just failed on it.

GUNTHER : That's a very interesting area, that whole computer-human interface area is a very difficult one to make any hard and fast judgments about, because there's no real easy way to analyze it. You can't analyze it in terms of a normal engineering perspective. It's something going on between the digital computer and this neural computer we have above our shoulders here. And you're trying to get these two computers to connect with one another, and it's not obvious how to do that. The Windows interface has now become accepted, but you and I know that that's not the last word in computer-human interface, we hope. So ultimately, this is going to go completely somewhere else. The Windows interface is the thing that's accepted right now. And not obvious even how to continue to improve that right now, it's very difficult.

DVORAK : Exactly. And not only that, it's not only not obvious, but the predicted interface of the future probably is wrong, and what they're starting to assume now is that some sort of combination Windows-voice, where you'll be talking to the computer is going to be it because they've watched too many Star Trek episodes.

GUNTHER : Yeah. Well, I think when people start talking about having some kind of probes that connect to your skin, and it actually picks up on brain waves, then maybe that's going to be something like that next real important interface.

DVORAK : Well, IBM has an interesting interface that's kind of... I don't know if it's an interface, but they have this thing where they scan your retina, where they have these beams. And then whatever you're looking at, they figure it's something you're interested in, and it goes off and does something, which I think is making one too many assumptions, usually, because, you know, a moment of attention doesn't necessarily

mean interest. But there's stuff like that, I guess, people are working on. But it seems that there has to be... you know, the labs like PARC, that you worked at, where these ideas were, you know, rampant, is where something has to come... something has to come from like some obscure research being done somewhere.

GUNTHER : Well, you know, I'm inclined to think also that that lab model, actually, I think is now outmoded. You know, that lab model really goes back to Thomas Edison, if I think about my physics background. That whole lab model–IBM Watson, Xerox PARC, Bell Labs–those are three major labs in the country in the last two decades or so; those labs are on their way out. The Edison model is failing because in a certain sense, all those labs were lucky... in fact, Xerox PARC really had quite incredible luck. I mean, not that these people weren't skilled, but also an incredible combination of luck to produce the Alto, basically, the Macintosh, in 1974, so bit map displays, mice, ethernet, laser printing–all of that coming out in that same decade or so it's an amazing piece of luck. And Bell Labs had luck with the transistor and so on and so forth, but none of them made a big commercial success out of it. But I tend to think now that the next line of innovation in this coming century is going to come from either individuals or small groups, which I think possibly is why this open development model that Netscape has now taken on recently, you know, where they open up the source code... I think that's the very, early beginnings of something along those lines. Because as an independent consultant, what I find is that often what I'm doing on my own time or on a client's time is actually a form of research. It's not recognized as such because I don't do it in a formal scientific way, and I don't publish in any scientific journals anymore. The reason is because it doesn't do me any good because I'm only talking to those people... like if I did an mathematical paper, I'm just talking to another bunch of mathematicians, and that doesn't do my business any good. So I don't write papers for those journals anymore, it's not worth my time. That doesn't mean I'm not doing research. And, in fact, I think this is... sometimes, you know, when I have a heady moment, I think this is the new form of research in the 21st century–small, localized groups, including individuals, working together across the Web, and so on and so forth. And this open source development, I think, captures another aspect of that same kind of

computer research in the future. I don't think it's going to come out of big labs anymore, I think all the big labs are going to disappear, as they, in fact, are.

DVORAK : I guess that makes me a researcher in some odd way.

GUNTHER : Absolutely. And you're communicating with a much broader base than most researchers.

DVORAK : Triggering other ideas.

GUNTHER : Yes, you're a catalyst.

DVORAK : Well, that's one way of putting it. Well, that's a very interesting theory, I think. And I don't see why it wouldn't be that way because everything's been more networked as time has gone by.

GUNTHER : That right. Just look a what's happening with the labs, they're doing the most craziest thing possible. Those labs were established to do research, and in many ways they have done so...

DVORAK : So that would mean Microsoft, with all its assertions about being the next big research company is essentially making another blunder here.

GUNTHER : Oh, I think that's the wrong paradigm, because all those labs that are supposed to have been doing research, are now... they are now positioning themselves according to corporate edict... corporate edict is Ye shall make product.' And that's totally crazy, because you've got people who have been trained to do research and they're good at it, and yet, they're going around saying, Well, I'm really working on product.' They know nothing about product. That's one of Xerox's major failings is to get product translated out of the lab into the commercial marketplace. So those people don't know how to do that, nor should they. That's not what they're good at. But they're going around saying, We're working on product now.' Same at Bell Labs, Lucent, IBM Watson, the same at Xerox PARC, that's what's happening now. So that's an erosion of the original research model Edison established 100 years ago or so. So I think it's just waning on that basis, and I just look at the empirical evidence.

DVORAK : But what about research in very high ends... I don't know if you want to call them topics or technologies, like semiconductor manufacturing or research in new types of silicon formulation? I mean, this can't possibly be done on an individual basis. You need a lot of cash behind you, you need a lot of equipment, you have to have a lab.

GUNTHER : Well, I agree with you, that's the perception. I'm not so sure about that. And the reason I say that is because about 10 years ago, I was involved with some people who had left PARC, and what they were trying to do was establish a chip foundry in a box. Now, just think about that model for a moment. We have things like, you know, a video recording, and audio recording on a PC. These things are called, you know, audio in a box' or video in a box.' Imagine if you could make a chip in a box. And they actually had the ideas for doing it. It didn't become a commercial success, I think, for business reasons not to do with the technology itself. So it's completely in mind, it's not out of the question, that a small group of individuals could do things like build chips. In fact, that's what students do. They use the Moses fab line down at ISI now to run their actual chips. You don't have to go through Intel. I do agree with you that that is the way things are going in terms of the microprocessor commodity marketplace, but that's not research.

DVORAK : And I guess another thing that makes your point is the development of RISC technology was largely done at a university.

GUNTHER : Yes, I would be rewinding a little bit on that argument, I would even be critical of universities. I think universities are in serious trouble as hot spots of research in the future. I think they are also being eroded in the same way that industrial labs are–it's the wrong model. A lot of my clients do not look to universities for technological innovation. They will pay me to do the innovation for them, to the degree that I can do it, or go and buy third-party product, or wait until something comes into place. I mean, they're more likely to support some form of research than necessarily go looking to universities for the answers. And if you look how the Web is impacting the way education is going, universities are almost no longer the seats of learning anymore because you can get graduate level degrees on the Web. University of Phoenix is one example, here in the Silicon Valley, that does that.

DVORAK : University of Phoenix is an interesting phenomenon, since it's really like a corporate university that you can get degrees from all over the country. It's very slick too. You know, in the old days–and I don't want to digress here, but I will–in the old days, you had you know, these mail order schools were always kind of second rate, sleazy-looking, you know, even their logo...

GUNTHER : We used to say you could get your degree in a Wheaties package.

DVORAK : Right. And it always looked like... you know, it didn't look very cool. These guys are very slick. I mean, you look like you've actually... you know, I mean, it's modern marketing approach to an education.

GUNTHER : Yeah, absolutely. And I'm not proposing that that's the actual representative model, I'm saying that's one thread. And the major point here is that the whole thing is being diversified, distributed. You can get your information on the Web, you can go to experts on the Web, you do not have to go to a campus to enroll, like you had to do in the past. And I think the old university model... the way I look at this is the current university model to me is like the cloistered academics of the dark ages, because there you had to be a member of the brethren, you know, the right religious sect and so and so forth, in order to be able to open up the books to get the information, otherwise, you're excluded, you go work in the fields. Now, we have universities, I think, that are getting themselves in that same position because they're trying to keep information localized in the cloisters, while we have the Web surrounding them everyday. Do you see what I'm getting at?

DVORAK : Oh, yeah, yeah, it's pretty obvious. I like it. We're talking to Dr. Neil GUNTHER, who is the head of the Performance Dynamics Company, and he's basically telling us that we're all doomed...

GUNTHER : Unless you hire me.

DVORAK : Yeah. You can find him in Mountain View, if you're looking for that. He's at `www.perfdynamics.com`. Anyway, Dr.

GUNTHER , thanks for being with us today.

GUNTHER : You're welcome.

Maureen MCGINLEY : You've been listening to Real Computing with John C. Dvorak, the insider's guide to personal computing. Real Computing is independently produced and receives funding from companies like Xerox, with Pagis Pro 2.0, a scanning software suite that offers full-featured photo editing and file management, so you can edit, organize, and share photos and documents via fax, e-mail, and the Web. Built-in Text Bridge Pro OCR eliminates the need to retype and reformat documents. Users can organize receipts, make color copies, and fill in forms on the PC. Pagis Pro is available at major retailers. Information is available at www.scansoft/realcomputing/.

And by Easy CD Creator, software that allows users to save data, music, photos and video onto a CD, then play on a CD ROM drive or an audio CD player. For Mac users, there's Adaptec Toast. Information about the latest version of Easy CD Creator is on the Web at www.adaptec.com/cdrec. This product is designed to assist in reproducing materials in which users own the copyright. See Adaptec's Web page for full legal information.

DVORAK : Well, I see by the digital clock on the wall that our time is up for this week's show. On next week's show, new technologies that will affect the way the Worldwide Web works in the 21st century. Until then, I'm John C. Dvorak. This is Real Computing, and everything you've learned this week will be null and void, by the way, by this time next week. So long!

MCGINLEY : Join us again next week for another edition of Real Computing with John C. Dvorak, the insider's guide to personal computing. Real Computing's mailing address is Real Computing, 1040 Greenwich Street, San Francisco, California, 94133. Or you can leave us a message at the Real Computing website, where you'll find the transcripts and Net Show audio files of the Real Computing program. The address is `www.realcomputing.com`.