The Lazy Person's Computing Crystal Ball

Last issue we completed our one-year, four-issue journey through a history of personal computing. Having stowed the Way Back Machine in the closet, it is now time to examine some new technologies and trends that I believe will affect how we compute over the next decade or so. Predicting the future can be fun as long as you don't bet large sums of money solely on blind optimism. My intent with this article is to give you some insight into what is actually in the works right now that will translate into real computing power in the near future.

Part of the fun with predicting future technologies is the *churn rate* at which they evolve. Intel cofounder Gordon Moore stated in 1964 that the amount of information storable on a given amount of silicon had roughly doubled every year since the technology was invented. Known as *Moore's Law*, this held true until the late 1970s, at which point the doubling period slowed to 18 months, where it still sits today. Thanks to this effect, there is an excellent chance that the computer equipment inside whatever box you are opening today has already been surpassed by something currently rolling off an assembly line. Regardless of how fast our technology evolves, though, I believe there are four things you can count on over the next 10 years: (1) Miniaturization. Devices will continue to get smaller. (2) Convergence. Devices will try to incorporate more functions. (3) Connectivity. The world will become more connected, as both wired and wireless networking advance. (4) Convenience. Technologies that make things simpler will win. At the end of this article I will also discuss technologies and trends that I think will die off in the next decade.

Twin Triumphs

Speaking of deadly technology, I have not given a detailed report on Zippy and his family since the twins were born. As you can imagine, life has gone through a few changes at *Chez Zip*. The twins, Cassie and Paul, are now a little over a year old. They are cute, cuddly, adorable little toddlers who love music, books and toys. Of course, their toys are a little different. Most one-year-olds don't have Palm Tungsten PDAs. It's just a little unnerving watching children who are still pre-verbal solving PDA games that beat me like a drum.

The big event last weekend was Zippy inviting me to his cabin in the woods. Despite his ineptitude with technology, Zippy has many other talents. He is a competent woodsman, fisherman and hunter, though his ability may be due more to luck than skill. The last time I went fishing with him, I watched him cast a large lure with three huge treble hooks about 100 feet from the boat and hit a six-pound bass on its head as the unlucky fish broke the surface. The lure stunned the bass and two of the hooks latched onto the side of the hapless fish.

Deer hunting for him apparently involves walking about 100 yards from his cabin, waiting five minutes, and then shooting whatever 8-point buck happens to wander by. He is living proof that it is better to be lucky than smart. His luck was not going to hold on this trip, though. After he got me in the car and we had driven too far to turn back, he casually mentioned he had made some "improvements" to his cabin. "Uh, what kind of improvements?" I asked, desperately hoping that they were anything but electronic. "Oh, the whole cabin is wired now," he replied with a wild gleam in his eye. It was too much to hope that the wires were connected to explosives, which would be less dangerous than letting him play with computers. I had one faint hope. "Uh, Zip, does your wife know about this?" "Oh, no! I wanted to surprise her when we go up there for vacation next month, so it's *Top Secret*. Promise you won't tell?" Since it was likely I would not survive the weekend in Zippy's "smart cabin," it would be an easy promise to keep.

A few years ago, a major power utility, television cable company, and software company, among others, started a project to create a "smart house." The goal was that everything in the home, including power, telecommunications, appliances, security, and all other electronic components could be controlled and monitored by the computer system. The whole system would be run by a universal remote with a friendly user interface, voice commands or automatic sensors. Zippy had volunteered to be a beta tester. I had a bad feeling about this.

Everything started out pretty well. Zippy let me drive while he used his cellular telephone to turn on the heat and lights, set the thermostat up few degrees, and preheat the oven so we could make pizza when we arrived. Much to my surprise, everything worked perfectly. Maybe this would work after all. I was impressed with the universal remote, a Palm-type PDA with wireless networking. Whoever had designed the interface knew what they were doing. Despite the small screen, everything was laid out logically and easy to find and activate.

Paradise, however, was quickly lost. The first problem appeared that night. The house has two very sophisticated systems that do not coexist very well: the entertainment and security systems. Zippy has a home entertainment system that would put most commercial digital movie theaters to shame. Remember the scene from *Jurassic Park*, where the impact tremors generated by the Tyrannosaurus Rex's footsteps appear in Jeff Goldblum's cup? With Zippy's sound system, they appear in your own drink cup, as well.

Unfortunately, playing the speakers louder than 25 decibels triggers the vibration sensors built into the windows. Imagine our surprise when the local police broke down the door midway through the movie because the security alarm had interpreted the window vibrations as someone breaking in and silently called the police. And, because the sensors were going off repeatedly as the movie progressed, the messages to the police had escalated from "*possible break-in*" to "*full-scale attack*" warnings. With apologies to the local constabulary, we turned off the security system.

The next morning the kitchen "crashed" when Zippy opened the refrigerator door to get his orange juice and the light bulb blew. A second later, everything else electrical in the kitchen went dead — refrigerator, lights, toaster, coffee maker, microwave, wall clock — everything. We unplugged and re-plugged all the appliances and reset the circuit breakers, but nothing would coax the kitchen out of its coma.

We called the Help Desk, where a nice lady ran some remote diagnostics. She said the problem was being caused by an "unanticipated failure mode." The designers had never anticipated having a refrigerator bulb fail while the door was open. When the bulb blew, the kitchen system interpreted the burn out as a power surge, went into "prevent/protect" mode, and shut down the entire kitchen. But when the electrical power system's sensors reported that there had not actually been a power surge, this set up a logic loop that confused the kitchen system, which then refused to respond to restart commands. The Help Desk lady swore that this was the first time this had ever happened. Restoring the kitchen took two hours.

The final straw was the virus that got in via the broadband connection after we had turned of the security system the first night. It seems the security system was a one-size-fits-all deal that protected everything from the windows to the network server and the climate control system. The security documentation didn't mention that turning off the window sensors would also shut down the firewall. Haven't these people heard of *segmented security*? When we woke up on the third day, the bedrooms were saunas, the kitchen was a skating rink, the freezer had defrosted, the washing machine was flooding the mud room, the water heater was venting steam and the stereo was playing the banjo music from "Deliverance." Lights were flickering like the strobes in a disco until they burst into shards from the strain. Flashing across the screen of every computer in the house were the words: "Be it ever so humble, there's no virus like HomeWrecker!"

We evacuated. The tech support team eventually had to disable the cabin with an electromagnetic pulse weapon. Zippy and I then spent the next few days having a wonderful holiday canoeing and backpacking in the forest, sleeping in tents and cooking over an open fire. Mother Nature did her best to cheer him up, but no matter how many fish he hit in the head with his lures, (not to mention the 12-point buck that dropped dead of an apparent heart attack not 50 feet from our camp), Zippy remained depressed and disconsolate. His dreams of creating a computer-controlled "Nerdvana" in the wilderness were dead.

There are a couple of lessons here as we move into using new technologies ourselves. First, trying to do everything all at once can be dangerous if you inadvertently mix technologies that compete for the same space. Second, no matter how much we plan for, we cannot plan for the unforeseen. We can only set up systems and processes that we hope will allow us to recover as quickly as possible when something catastrophic happens. Finally, new is not always better, particularly if it makes us so dependent that we no longer understand how it works. No matter how sophisticated a system is, someone still needs to be 10 percent smarter or we will abdicate so much control that we cannot take charge when we need to. With that cautionary tale under our thinking caps, let us now turn to some of the technologies that I expect will influence computing over the next 10 years.

Size Matters

Earlier in the article I predicted that devices will continue to get smaller and that we will see more and more multifunction devices. This is going to require some leaps in technology to reduce the size of computing components. There are three components that determine the size of our current computing devices: storage media, processors and display area. The first two are affected by physical limitations of what you can pack into a square inch, and the third by what you can reliably produce given the available viewing area. We will begin with storage and processing.

Nanotechnology is the science of making devices with features measuring less than 100 nanometers. A single nanometer measures roughly 10,000 times smaller than the width of an average human hair. Nanotech is not new, though. In the early 1990s, futurists predicted a booming market for microelectromechanical systems (MEMS). However, difficulties in manufacturing, packaging and other problems with producing MEMS so small that they have to be viewed with high-powered microscopes slowed development.

The most prominent effect of nanotech in the computing world will be on the size of storage devices. A recent report by NanoMagnetics, a British company, claims that computer hard drive capacity could be increased a hundredfold by using a common protein to fabricate nano-scale magnetic particles. Their process uses a common protein called apoferritin to create a material consisting of magnetic particles each just a few nanometers in diameter. Apoferritin is, by the way, the main molecule in which iron is stored in the human body. Each particle can store a bit of information and together they can be packed onto a disk drive at much greater density than is possible using existing hard disk manufacturing methods.

Seagate, a well-known manufacturer of computer storage devices, is also reportedly working on a nano-magnetic material that is fabricated chemically, rather than by using proteins. At the moment, we can pack, at best, 450 gigabits of data onto a square centimeter of conventional storage media. With nano-storage, this could eventually be improved to anywhere from 3,000 to 5,000 gigabits per square centimeter. Combine that with other predictions that new computers in 2005 will come with 500 gigabyte hard drives, and you can see where this is all going. In apparent violation of various physical laws, computers will get smaller on the outside and bigger (virtually) on the inside. The companies working on this technology believe that it will surpass conventional hard disk density by the end of this year and that a nanotech-based storage material could be available between 6 to 10 years from now. That means the question isn't whether we will get storage based on nanotech, but when.

Another application of nanotech will be in chips. NASA scientists have reportedly developed a new manufacturing process that uses extremely tiny carbon nanotubes instead of copper conductors to interconnect parts within integrated circuits. The main advantage of using carbon nanotube interconnects within integrated circuits is that they will have the ability to conduct very high currents, reportedly more than a million amperes of current in a one square centimeter area. Because copper's resistance to electrical flow increases greatly as the metal's dimensions de-

crease, there is a limit to how small copper conductors can be. However, because carbon nanotube electrical resistance is very low, they will allow development of smaller chips that use less power and produce less heat, allowing manufacturers to increase the number of circuits on a chip beyond the current limits of copper circuits.

As chips and storage get smaller, they will subsequently allow much smaller "boxes" for computing devices. Fifty years ago, a single computer would fill a large room. Today, there are handheld calculators with more processing power. In ten years, your watch may outperform today's 3 GHz Pentiums. As with any new technology, it will take time to absorb nanotech into computing. After all, it took 15 years from the invention of the transistor in the 1940s before it became a real player in the electronics industry. However, there is one other limiting factor concerning miniature devices, which will become the principal constraint on miniaturization: display size.

Electronic Paper

One of the ergonomic problems with using computers is that the screens cannot approach the clarity of ink on paper. The cathode-ray-tube (CRT) screens on most desktop PCs achieve a resolution of 80 to 100 pixels per inch (ppi), at best. The human eye, on the other hand, can reliably distinguish around 200 ppi. Good quality liquid crystal displays (LCDs), now seen on what most people recognize as "thin, flat-screen" monitors, have better resolution than CRTs. The best commercial LCDs resolve up to 120 ppi. While this is better than CRTs, it still is not as easy to read as paper and ink. By comparison, images in most commercial magazines are printed at around 150 lines to the inch, better than the best LCD, but you can still see the dots if you look hard.

However, better screens are on the way. IBM Research has reportedly developed an active matrix LCD named "Roentgen" that has a resolution of 200 ppi, or 40,000 pixels per square inch. For normal viewing, that is pretty much the same detail our eyes can see on paper. In addition to Roentgen, various other researchers are working on tiny, lightweight and low-power virtual displays for portable electronic devices like pagers, cellular telephones and wearable computers. Current color cell phone screens display 320x320 pixels, with large displays running at 320x480 pixels. Considering the amount of information the average Web site tries to push that does not give you a lot of room to work with. There are still low-resolution screens where you can count every pixel. Any advance in resolution will theoretically be a good thing.

Toy Shopping

So what are we seeing now for new convenient, compact or multifunction devices? Well, first up would be the WRISTOMO "wrist phone" being built by Seiko for the Japanese market. As its name implies, it is a telephone worn like a wristwatch. It allegedly supports Web browsing at up to 64 kbps and can receive and transmit e-mails with maximum size of 3,000 characters. Weighing in at 171.5x40.4x18.5mm and 113g, it is advertised as being able to provide 120 minutes continuous talk time.

A bit closer to home are combination devices like the Handspring Treo 300 or the Kyocera 7135 cellular telephone, both of which include a fully functional Palm PDA. In addition to integrating the hardware, these new combination telephones are integrating their functions. You can share one contact database between your computer, PDA and cell phone. You can tell your telephone to "call Russ Fraser," and it will pull the name from the PDA contact database that came from your PC and then dial the number.

Another technology integration we will see more of is "voice over Internet Protocol" (VoIP). Until 20 years ago, telephone networks were primarily analog systems. Over the last two decades, though, digital telephone technology has become more prevalent. Most of the major players in the voice telecommunications industry are pushing to move voice communications onto IP networks. It is not a completely mature technology yet, but it does offer certain advantages over traditional telephone systems. I plan to address VoIP more thoroughly in a future article.

I'll Meet You on the Network

Before we close, here is one more useful new technology to look at: self-configuring networks. Last year, Apple Computer demonstrated a technology called "Rendezvous." It is apparently the first mass-market implementation of what can be called "zero-configuration networking," and allows devices to talk to each other without requiring manual configuration. Apple first demonstrated Rendezvous using their iTunes music management software. It demonstrated that a user holding a Rendezvous-enabled laptop with a wireless networking card could walk into a room and automatically see the iTunes music files of everyone else in the room with similar systems. That is cool if you like file swapping, but what I am really looking for is something business-oriented that won't get me raided by a Recording Industry Association of America SWAT team.

Apple is delivering that now. Over the last year, Apple and some associated vendors have embedded Rendezvous in an increasing number of applications. Most of the major printer manufacturers that support Apple systems have upgraded their machines to support Rendezvous. One that network administrators might appreciate is that Apple's new Safari browser allows you to change Rendezvous-compatible printer configurations without having to hunt down specific IP addresses. The practical impact of this is that if you are far away from your office and need to print a document from your wireless laptop, you can do it from any Rendezvous-capable printer within range of your laptop. You do not even need to be on the office local area network or logged into any directory software.

What would it be worth to no longer need to configure a computer for network printing? This routine normally involves wading through dozens of folders in search of the proper IP addresses for your office printers, and I defy anyone other than a hardcore techie to write, from memory, the path and name of even one of the printers currently set up on their computer. Most networks today require file and print servers. With Rendezvous, you won't need them. With the functionality that Rendezvous provides, we could reduce the cost of managing the network as a whole by using Rendezvous-ready software to broadcast printer and file sharing service changes to every machine on the network, which could eliminate the overhead and system bugs inherent in individually changing settings on each desktop.

On the storage side, network storage device maker Chaparral has built Rendezvous into the latest version of their storage-management software. Configuring a network that uses Chaparral storage is now allegedly a point, click, configure and go have coffee

process. You can allegedly back up the contents of the network to the storage device with pretty much zero tweaking. Finally, Rendezvous is finding a role in database management. Sybase has apparently built Rendezvous into its client software, which allows authorized machines to log into Sybase databases without any additional configuration. How many database administrators would just like to authorize new users without reconfiguring their machines, as well?

Individually, each of these little tweaks might not count for much. But if you add enough of these functions together, it becomes clear to me that it will cost less to run a network with Rendezvous-like functionality than it does to run today's networks. One last thought: It is rumored that Rendezvous will have the ability to check processing usage on other Rendezvous-enabled machines around the office and share processor-intensive tasks. Distributed parallel computing is useful; having it without configuration issues would be cutting edge. As usual, I don't expect everyone to go out and buy Macintoshes just to get Rendezvous. Eventually someone will migrate the technology to Windows or Unix. My point is that technology like Rendezvous, which is firmly grounded in *convenience*, will become a force as our networks evolve over the next 10 years.

Fond Farewell

I would like to close with my list of things that I hope either burn out or fade away. Here they are:

1. Lock-In Licensing. "*Agree to send me money forever now and save; buy later and we will be forced to charge you more.*" As I predicted last spring, licensing and support issues are starting to make open source software more attractive this year. The companies with restrictive licenses are making modifications almost monthly to try and attract or keep business. But there are some who have moved on and are not looking back.

2. Proprietary file formats. While I understand that the makers of word processing, presentation and database software have a desire to protect their market share, proprietary file formats remain one of the biggest barriers to progress in the computing world. Web staples like Hypertext Markup Language (HTML) and Extensible Markup Language (XML), and open source software are making inroads here. People are moving away from proprietary solutions and toward competition on merit, not inertia. Again, companies that own proprietary standards are opening them up to compete.

3. Spam. Junk e-mail allegedly now constitutes almost half of all Internet e-mail traffic. Laws against it have been ineffective largely because of jurisdictional issues. However, I remain hopeful that some combination of legislation, international treaty, and ISPs choking off spammers' air supplies will resolve this before spam chokes the Internet to death first.

That's all for this issue. We will come back to this in a few years and see how I did. In the meantime...

Happy Networking!

Long is a retired Air Force communications officer who has written for CHIPS since 1993. He holds a Master of Science degree in Information Resource Management from the Air Force Institute of Technology. He is currently serving as a Telecommunications Manager in the U.S. Department of Homeland Security.

Cont'd from page 35

as well as communications and gateway nodes to perform such missions as signal intelligence collection and radio jamming — performing whatever RF mission is required to deny enemy access to the RF spectrum.

JTRS will eventually be integrated into 64 different types of aircraft in the Air Force inventory, as well as a myriad of ground-based platforms, during scheduled depot level maintenance. JTRS is programmed to replace 750,000 radios within the inventory covering operations such as navigation, positioning, location, identification, Air-to-Ground, Air-to-Air, Ground-to-Ground and satellite communications. Capitalization of integration costs will be realized by reducing 124 different radio sets to approximately 10 to 20 form-fit radio sets.

There are several challenges yet to be resolved to fully exploit the inherent capabilities that JTRS will bring to the battlefield. Two of these include antenna research and legal constraints affecting all radio systems, especially those designed to access the entire RF spectrum, due to potential conflicts with non-U.S. authorized frequencies and non-military systems. While several challenges have yet to be resolved to fully exploit the capabilities that JTRS will bring to the battlefield, JTRS is a revolutionary way of doing business. Once fielded, warfighters will no longer think of the RF spectrum in terms of hardware but as capabilities. Multiple software modules will allow implementation of different standards in the same radio system (including the capability to employ multiple waveforms resident on the same set). Radio receivers will be reconfigured over-the-air, thereby reducing maintenance requirements.

"In the past, if you were equipped with an HF radio, you were limited to communicate with HF waveform subscribers. Once JTRS has been fielded, a warfighter will be able to talk to another warfighter on multiple waveforms, and it will be totally transparent to him that this is what he is doing. We are only beginning to appreciate the realm of the possible that JTRS brings to the battlefield," concluded Col. Whitehurst.

Capt. Todd White is the Chief of Public Affairs for the Air Force Command and Control and Intelligence, Surveillance & Reconnaissance Center at Langley Air Force Base, Va. Commanded by Maj. Gen. Robert F. Behler, the AFC2ISRC is the lead organization to integrate and influence command and control, intelligence, surveillance and reconnaissance strategies, roadmaps, and investment plans for the Air Force. The center also oversees the Air Force Experimentation Office (AFEO) in Hampton, Va.; and the Command and Control (C2) Battlelab, located at Hurlburt Field, Fla. In addition, the center has Officer Liaisons (OL's) assigned to 16 additional agencies.

Our sincerest apologies to Mr. Cray Henry for incorrectly listing his name in The High Performance Computing Modernization Program article in the Spring 2003 issue of CHIPS. Our thanks to Mr. Henry for his graciousness. Please visit the HPCMP Web site at: www.hpcmo.hpc.mil.