Columns 2005 to 2007

Peter Rysavy wrote columns for the Network Computing Mobile Observer from 2005 to 2007. He has also wrote columns for Unstrung and Mobility Loop.

Index:

Sprint Nextel's Jump to the Future, Feb 8, 2006
Mobile Content - Getting Better Slowly, Feb 3, 2006
The Resurrection of IEEE 802.20, Jan 25, 2006
The Smartphone Conundrum, Jan 18, 2006
Mobile Phone Music, Jan 12, 2006
The Perils of the Bleeding Edge, Jan 5, 2006
HSDPA Speed: Lab versus Real World, Jan 2, 2006
Municipal and Mesh Wi-Fi, Dec 21, 2005
Ultra Wideband Standardization, Dec 18, 2005
Mobile Broadcast Technology Developments, Dec 7, 2005
Linksys WRT54G--Update, Dec 7, 2005
Linksys WRT54G--Not What It Used to Be, Nov 30, 2005
EV-DO Revision B: Speed Hype, Nov 23, 2005
Mobile Commerce Rebound, Nov 18, 2005
Confusion About 4G, Nov 9, 2005
WiMAX World, Nov 3, 2005
Wireless Data Pricing--One Hundred Times Less Expensive and Still Not Satisfied, Oct 19, 2005
Embedded Wireless WAN, Sep 28, 2005
Redefining the Endpoint--Wireless Broadband Routers, Sep 7, 2005
High Speed Downlink Packet Access (HSDPA), Aug17, 2005
The Uneven Uptake of Wireless Data, Jul 27, 2005
Where is WiMAX?, Jul 6, 2005
Wireless E-Mail for the Masses, Jul 15, 2005
Parallel Mobile Universes, May 25, 2005
Beyond the Airlink - IP Multimedia Subsystem, May 4, 2005
3 Sides to 3G, Apr 13, 2005

Handheld Computing Evolution

Network Computing Mobile Observer, June 13, 2007

We think of our mobile telephones as phones or PDAs, and sometimes as music players (including the Apple iPhone that launches this month) or cameras, but let's be clear: These are computers. We just don't think of them as general-purpose computers, since not even the smartphones are as "general" as our laptops and desktops. However, there are huge implications in the increasing computing capabilities within handheld form factors, combined with broadband capability over Wi-Fi and 3G networks and huge storage capabilities. I've written on this topic in the past, but a number of new developments warrant further discussion. First, I've just completed a large private handset evolution study working with Datacomm Research, where we interviewed 44 organizations

that provide all the various components and capabilities of mobile handsets, including operating systems, chipsets, Wi-Fi, ultra wideband, 3G, user input, storage, power, fixed mobile convergence and so on. Second, there have been some recent industry developments--including those by Apple and Palm--that portend interesting trends.

The one by Apple was CEO Steve Jobs recently indicating that the company is putting serious thought into how to allow third-party application development for the iPhone. You'd figure that Apple, with all its innovation in computing, would be well positioned to examine the handheld market with a fresh perspective. So here is Jobs, realizing that the company's iPhone, which runs Apple's Unix-based OS, potentially will be able to run a vast number of applications. But at the same time, he realizes that for the platform to be successful, it has to provide a superb user experience with respect to stability, ease of use and enjoyment--something sadly lacking in many smartphone systems today.

What's of interest is that there isn't a clear-cut answer. Do you target devices like this with a fixed set of smoothly interacting applications, or do you open up the system as more of a general-purpose platform? My view is that a prepackaged set of applications may better address the consumer market in the short term, but that over time people will expect a range of applications that can only be supported through third-party development.

One of the key conclusions of our handset study was that while general-purpose smartphone operating systems like Symbian and Windows Mobile (and soon mobile Linux) are the base for today's high-end smartphones, they will increasingly power the mid-tier phones of tomorrow. There are a number of reasons for this. One is that the hardware costs to support these more powerful OSs are going down quickly. Another is that as the applications become more sophisticated, more complex and more diverse, the ecosystem of software component suppliers will get larger, simply because no one company will have the engineering resources to do it all.

For example, sooner or later our handhelds will need security defenses as thorough as those on our desktops, meaning that solutions will need to come from companies that specialize in security, such as Symantec. And that can only happen in systems that have well-defined interfaces, and it can only happen for a handful of platforms. It's no different than the forces that have resulted in only a small number of operating

systems on our desktops and our servers. As more applications are crammed into phones, the proprietary real-time operating systems that power most phones today simply won't be able to keep up, except in lower tiers of the market.

Greater handheld capabilities raises the following question: If handheld devices have all the computing power users need to run applications that matter to them, why would they want to maintain a separate laptop computer? The answer is an extension of the trend of people having separate desktops and laptops when laptops first came out to the now common use of laptops for all of people's computing needs. The "handheld for everything" approach is exactly what Palm just announced with its Palm Foleo. Palm describes the Foleo as a mobile companion device with a 10-inch display and full-size keyboard that is designed to work with smartphones. It turns on instantly, connects to the phone by Bluetooth, has Wi-Fi and is expected to cost $500 when available this summer.

I think that we're a little early in the game for this kind of solution, simply because the number of available applications with this particular system is quite limited compared to what's available for laptops, which don't cost much more. However, I do think it's a sign of things to come, namely that the handheld has all the computing, storage, communications and entertainment users need, and all that's missing is a larger display and keyboard for some of the applications. So kudos to Palm for being the first to suggest this approach, especially since it addresses the obnoxious problem of having to maintain multiple computing environments--no easy task as smartphones get increasingly more complicated.

Finally, there is the ultra-mobile PC trend, another example of the potential for miniaturizing computers. The issue here is that user interfaces like Microsoft Windows don't scale well to tiny displays.

Where all this is heading is not completely clear, but I do think one inevitable consequence will be more powerful general-purpose operating systems in an increasing percentage of phones.

Global Mobile

Network Computing Mobile Observer, May 23, 2007

I recently completed a combined business and personal trip to China,

spending time in Beijing, Shanghai and nearby areas. To maximize my connection opportunities, I brought a Cingular/AT&T 8525 that supports quad-band GSM/EDGE (800 MHz, 900 MHz, 1.8 GHz, 1.9 GHz), tri-band UMTS/HSDPA (800 MHz, 1.9 GHz, 2.1 GHz) and Wi-Fi with IEEE 802.11g. To expand communications options further, I installed Skype on the phone, with the idea of using it over Wi-Fi. I figured if this combination would not let me stay connected, nothing would. I also took a close look at pricing for international usage, so I wouldn't have any rude surprises the next time I received my bill. As it turned out, my ability to communicate was superb. But I had to closely watch my usage, as overseas usage without an international usage plan is extremely expensive.

The first stop was in Tokyo for a few hours before connecting to Beijing. I didn't have high expectations, as GSM is not available in Japan. However, my phone promptly registered onto the NTT DoCoMo network in UMTS 3G mode, and I was able to easily download my e-mail. That was exciting, even if the actual e-mails weren't.

The next stop was in Beijing. My phone registered onto China Mobile's GMS/EDGE network, and connectivity was a breeze for voice and data. Coverage was never an issue, with my phone almost never showing less than three out of four bars of signal strength. This was no surprise given all the cell towers I saw, both in the cities as well as in rural areas. As of the beginning of 2007, China Mobile had more subscribers than the population of the United States, so I suppose that takes a few cell sites to support. Even on a tour to the Great Wall I had constant connectivity, including walking on the Great Wall. This was impressive, given that we were in the mountains, albeit in a dense tourist zone. Data operation was a breeze, though occasionally I had a problem establishing the initial data connection. But once I had it, I could keep the data connection going all day with no issues, and I didn't run into this problem in other areas in China. That was the good news for connectivity.

The bad news was the expense. Without having signed up for an international plan, data roaming cost $19.50 per megabyte and voice usage was $2.29 per minute. For voice, this is clearly extremely expensive, but for data, it's not as bad as you might think. Granted, a typical hour of Web browsing that consumes 10 MB of data would have cost $200, but judicious use of e-mail with a 5-KB maximum message size actually translates to just 10 cents per e-mail. You could even reduce

the 5-KB value to 2 KB, but then you would get too many incomplete e-mails. And the only way of getting the rest of the e-mail with Windows Mobile Outlook is to download the entire rest of the message. The net of this careful e-mail usage was that I only spent about $30 on e-mail over 10 days. I didn't survey the other operators for their global pricing, but I would expect it to be similar.

I stayed in a Westernized hotel with good amenities that offered international dialing from the room, which I knew would be extremely expensive, Ethernet for Internet access and a business center that had PCs as well as Wi-Fi access to the Internet. At first I thought that using the business center's Wi-Fi would be less expensive than cellular data, but I never even bothered, as the cost was about $4 for a 15-minute minimum usage.

As for voice, I thought about Skype over the hotel's Wi-Fi network. My prior research showed that people successfully use their Skype Out accounts for global calling at 2 cents per minute, but as noted, the Wi-Fi cost was very high. I also tried Wi-Fi visiting some friends, but they had locked down their network with MAC address filtering and didn't know how to change it. Instead, I purchased what is called an IP phone card for China Unicom, and that allowed calls to the United States at about 30 cents per minute. For the limited number of calls I had to make, this was fine. The hotel surcharge for dialing the access number was just a few pennies per minute.

Though I went to China for this experiment, usage in the more than 200 countries supporting GSM would have been about the same for both connectivity and pricing. Four days after I got back, AT&T announced a global PDA plan with data access priced significantly lower, at $70 for 20 MB. There are also international BlackBerry plans. In addition, Sprint and Verizon have new international plans, though their CDMA networks require you to obtain dual-mode GSM/CDMA devices or to rent a GSM phone. These international plans make sense for regular overseas travelers but not for people traveling only occasionally.

Bottom line, getting wireless voice and data anywhere in the world has become extremely easy. Just be careful how you use it.

Wireless Networks, Wired Bottlenecks

Network Computing Mobile Observer, May 2, 2007

Are you feeling backhaul constrained? You may not know it, but that may be the reason you're not getting the throughput you're expecting with your wireless applications.

While there is plenty of press about advances in wireless technologies, especially about increasing speeds and lower latencies, it's actually likely that wireline technologies will provide the next boost for wireless application performance. That's because today, many networks are constrained by the backhaul from the wireless access point or cell site. If you've ever been at a wireless hotspot and done a throughput test, you've probably noticed a maximum rate just under 1.5 Mbps. That's because many wireless hotspots use a T1 circuit for backhaul, which has a maximum supported rate of 1.544 Mbps. So it doesn't really matter if you're using IEEE 802.11b or IEEE 802.11g, your throughput is going to only be as fast as the weakest link in the chain. Whether it's the 6 Mbps maximum over-the-air throughput of 802.11b after protocol overhead, or the 25 Mbps of maximum over-the-air throughput of 802.11g--both are constrained by the T1 backhaul. IEEE 802.11n will eventually increase range but it won't do anything for throughput.

It's no different in municipal Wi-Fi networks, many of which today provide 1 Mbps service, again due to backhaul constraints. The same situation exists at home, though cable modem and DSL speeds for many users are faster than T1. As an aside, I find this a little ironic given that a T1 used to be the gold standard for Internet connectivity, but is now much slower than consumer-oriented alternatives.

When we turn to cellular networks, it's the same story. Many cell sites have a couple of T1s going to them, which have to support voice and data traffic. If a cellular operator wants to provide, say 7.2 Mbps High Speed Data Packet Access (HSDPA) service, those bits may travel over the air at that blistering rate, but then grind down to backhaul-circuit throughput rates. My understanding is that this is not really a problem today because the number of 3G data users is still relatively small, and throughput is not advertised at these higher rates. Expectations today are more in the 1 Mbps range for peak 3G throughputs. But as the number of users ramps up, and as 3G has to start competing with other emerging technologies--such as WiMAX--operators will have to beef up the backhaul, possibly to four or five times the capacity they have now. That's a big jump, but I believe this is inevitable.

There is no shortage of market activity to address this trend. If the wireless site is already serviced by fiber, provisioning additional virtual T1s is easy. However, there are a lot of other optical transport approaches available, including metro Ethernet, Pseudowires (emulation of native services over packet switched networks), Multiprotocol Label Switching, as well as traditional SONET/SDH. There are also point-to-point radio options, including 60 GHz unlicensed radios that offer 100 Mbps throughput over a mile or so, and plenty of options below 60 GHz. In fact, this radio market is extremely competitive, with prices coming down quickly, and is likely to be the most attractive option for sites where fiber is not feasible.

Carriers are generally tight lipped about what they are doing, but I think they are all very active in their efforts to beef up both their backhaul and core networks to deal with much higher data flows. Assuming

wide-area wireless data finally takes off, solving this infrastructure problem is essential for both current 3G networks as the number of users increases, and especially for enhanced 3G offerings such as HSPA+ and EV-DO Rev B. Further out, systems such as WiMAX, 3GPP Long Term Evolution and Ultra Mobile Broadband will require even more backhaul capacity. Interestingly, it may ultimately be the backhaul and core network architecture that plays a bigger role in overall performance and operators with strong wireline networks may be at an advantage here.

There's lots happening on the home front as well. VDSL2 can provide up to 100 Mbps of theoretical throughput over copper wire, though only over fairly short distances. But 25 Mbps may be practical. Similarly, we are seeing cable modem rates approaching 10 Mbps, and fiber to the home offering 30 Mbps. Hopefully, before long, you'll be seeing the full capability of that Wi-Fi or 3G connection. But until then, make sure you understand where the bottlenecks are.

QoS and VoIP With 3G: Not an Easy Marriage

Network Computing Mobile Observer, April 11, 2007

There's lots of confusion out there about QoS, VoIP and 3G--these don't work quite the way many people think. This is an area I've been studying intently as it pertains to various projects, and after hearing a Qualcomm rep speak recently on the topic in a technical presentation on EV-DO Rev A, I thought I would clear up some misconceptions and organize my own thoughts.

Sprint and Verizon are upgrading their CDMA2000 EV-DO (Evolution-Data Optimized) networks from Rev 0 to Rev A, and one of the features is QoS (quality of service). QoS is also available in UMTS (Universal Mobile Telecommunications System) networks, including Cingular's HSDPA (High-Speed Downlink Packet Access) network. QoS raises two interesting questions: What applications need QoS, and what's the best way to make QoS available as a feature for customers?

A good place to start is VoIP. Many popular VoIP systems are based on best-efforts approaches only, and as a consequence, voice quality is hit-or-miss, as anybody who uses Skype extensively can attest to, especially with international use. The Internet at large does not support QoS protocols, although there are managed domains that do. This will change over time, especially as the backbone networks carry more time-sensitive voice or video traffic. Even in the absence of QoS mechanisms, wired networks have two huge advantages over wireless. The first is the ability to easily add more bandwidth. The second is consistent bandwidth with low error rates on any individual link.

Wireless networks, in comparison, are constrained by highly restrictive spectrum allocations and must cope with high link level error rates. To do so, they are highly adaptive, changing the amount of coding and modulation to provide best-available throughput at any moment. The result from a QoS point of view is problematic: a constantly varying amount of available bandwidth per user. So what are the options for 3G?

There are different ways to control QoS. Full QoS control means being able to control parameters such as throughput, delay, variation in delay (jitter), error rate and whether packets are sent in sequence. However, since QoS is ultimately a matter of packet prioritization, a much simpler QoS implementation would be to designate classes of traffic, where higher-class traffic gets through before lower-class traffic.

The first way that 3G will seriously use QoS is with 3G networks employing VoIP. Sure, you can run Skype over 3G today. But actually running VoIP with the same voice quality and capacity as current circuit-switched systems is complicated, which explains why nobody has done it quite yet. Circuit-switched methods are highly optimized, but most VoIP approaches are not. IP headers alone can kill you, with some 40 bytes of IP, UDP and RTP packet headers of overhead for every 22 bytes of VoIP payload. Techniques such as Robust Header Compression (RoHC) employed within the network bring this down to 4 bytes. If the cellular operator then applies full QoS for voice call admission and bandwidth management, and additionally applies a variety of other complicated low-level bit twiddling such as dejittering and time warping, the result is high-quality and high-capacity VoIP equaling and eventually surpassing current circuit-switched voice. This is on the road map for both E-DO and UMTS, though EV-DO operators are more highly motivated to make this work because their EV-DO radio channels currently support data only, whereas UMTS/HSDPA channels allow simultaneous voice and data. Note, however, that this application of QoS is purely within the operator's domain. It could, if the operator chose to do so, be made available to improve your Skype session. But I see this as unlikely in the near future.

I think it's likely that 3G operators will leverage QoS for their own applications, including VoIP, as well as for video services. But as far as exposing interfaces for customer applications to request parameters such as a certain amount of bandwidth or delay control, that's a harder sell. The problem is that with finite capacity, do you really want one set of customers using up all the bandwidth and thus denying data

service for others? And how much would customers be willing to pay for such a service? Plus, remember that the amount of radio resource required to deliver certain bandwidth depends on the users' instantaneous radio conditions. I know operators are scratching their heads over this one.

However, the class-of-traffic approach might be easier to monetize, as well as to manage. Here, the operator could provide higher-priority traffic for, say, emergency services or critical medical applications. Now, when all hell is breaking loose somewhere, the police and firemen using a public system could still be able to communicate. I'm not aware of any operator offering this yet, but it wouldn't be that difficult for them, at least technically, to do so.

Interestingly, some companies, not waiting for operators to expose QoS control, have taken matters into their own hands. For example, NetMotion Wireless with their Mobility mobile VPN, permits prioritization of traffic within its own tunnels, allowing VoIP, for instance, to get through ahead of other applications, and even blocking some applications, such as Windows Update, when on slower connections. Bottom line, QoS for wide-area wireless has sort of arrived, but it will be sometime before you see widespread usage.

Learning From Bluetooth

Network Computing Mobile Observer, March 21, 2007

Bluetooth has always had tremendous potential, though difficulty in using it has hampered adoption. However, if all goes as planned, personal-area networking will be much easier in the future.

Bluetooth is a wonderful technology, marred only by being extremely frustrating to configure. I have spent more hours setting up Bluetooth connections than any other kind of networking technology. However, once things are working, I've found Bluetooth extremely reliable and useful.
Here are a couple of my recent Bluetooth experiences, the issues involved, and why efforts underway should make life better with next-generation personal-area networks such as ultra wideband (UWB).

Most recently, I tried to configure a Cingular 8525 phone as a modem. And who wouldn't? The phone supports High Speed Downlink Packet Access (HSDPA), a 3G service that Cingular indicates has average throughput rates of 400 Kbps to 700 Kbps, and burst rates of over 1 Mbps. Moreover, the service allows simultaneous voice and data, meaning you can make or receive calls while on a data connection. I had an older Anycom Bluetooth PC Card in my laptop that I have successfully used for modem connections with the Bluetooth Dialup Networking (DUN) profile and a

variety of other Bluetooth phones in the past. To my dismay, I spent hours trying to get this card to work with my phone. With help from an engineering contact at Cingular (AT&T), we determined that the latest version of Windows Mobile 5 (AKU3) no longer supports the Bluetooth DUN profile. Instead, using the phone as a modem should be done using the Bluetooth PAN profile. After all these years of IT managers finally figuring out Bluetooth DUN, I'm sure Microsoft's phone is ringing off the hook about this decision.

I tried to get an explanation from an unofficial contact at the Windows Mobile group at Microsoft, but received no reply. Okay fine, I thought, I'll use that profile instead, but I had no luck. My Anycom card did not seem to support that profile. So I called Anycom technical support, but learned my product was no longer supported. Finally, with some intermediate steps that I don't have to space to recount, I purchased a Toshiba Bluetooth USB dongle (meaning the Bluetooth device plugs into a USB port). Based on prior experience, I know Toshiba makes good Bluetooth products. Sure enough, phone tethering with the mobile phone worked perfectly using the PAN profile. With Microsoft Windows seeing the phone as a network (instead of as a modem), connections were automatic and I must confess, more intuitive than using a dial-up model. For example, you don't have to enter strange dial numbers or obscure username/passwords. So Microsoft is right. Other than the confusion factor from people trying to use DUN, this looks like a better way to go. At this point, I had over four hours invested. I should mention that tethering using a USB cable[[isn't the dongle a wireless one? Where's the cable go??]] is much simpler.

The next step was to use Bluetooth for ActiveSync. Sure, I can use the cable, but sometimes it's nice to sync the device without needing a cable. Once again, my initial attempts here would not work. After an e-mail cry for help to Toshiba, I learned that this is challenging to configure, and they weren't sure what my specific problem was. I was only slightly discouraged. Finally I determined that Toshiba installs COM7 as a Bluetooth COM server on the computer, meaning that was the COM number I needed to enter into ActiveSync. I also needed to make COM9 an outgoing COM port on the phone (obvious, eh?). Then things worked perfectly. That was another four hours. Crazy, I thought. What normal human being would figure all this out. Let me note, however, that getting the Bluetooth headset working (the excellent Motorola H700) took only a few moments.

Anxious at this stage for some good news about personal-area networking, I called MCCI, a leading company that delivers USB drivers under the hood to the industry, and is actively involved with UWB. The next wave of PAN technology, UWB is exciting because it provides extremely high data throughputs, all the way to 1 Gbps. There is widespread industry support for two types of interfaces. One is Wireless USB, the other is WiNET. I'll discuss these in a future column, as there is lots of good stuff here. Bluetooth itself is evolving to support UWB at the lower layers, allowing easy migration of Bluetooth applications. MCCI informed me that there was more good news. Ease of use is a huge priority for UWB developers, especially in pairing devices together. The Bluetooth Special Interest Group is leading the charge here with something called Simple Pairing. One planned approach is to use Near Field Communications (NFC), another hot emerging wireless area, where

connections occur over extremely short ranges of just inches. With NFC simple pairing, you'll just touch two devices you want to communicate with each other and you'll be done. I can't wait!

Disconnect!

Network Computing Mobile Observer, February 28, 2007

Being always connected may be good for productivity, but it can wreak havoc with peace of mind.

I've spent my entire wireless technology career, now going into it's fifteenth year, working to promote wireless technologies and mobile computing, and especially their ability to deliver critical information to any location at any time. I've used wireless computing in almost form, beginning with Radiomail over the RAM Mobile Data network, and am forever experimenting as part of my consulting practice with a constant stream of new handhelds that I evaluate in various operating conditions. While I remain enthusiastic about the power of the technology, and believe that we're still in the infancy of how the technology will permeate and affect our lives, I think it's time to start coming to grips with the effects of always being connected on our psychological state.

I started thinking about this recently when I sent an urgent e-mail to somebody I was working with at another company. I received an auto reply saying "I am currently on vacation until Feb 28 and I will not be checking my e-mail." That was okay with me so I start making other arrangements. Yet, within 10 minutes this person responded to my e-mail. I guess his auto reply really meant "I am on vacation and I am getting my e-mail wirelessly. I will respond if it's urgent, otherwise I have a reasonable excuse to not reply at this time."

As a consultant, I can work from almost anywhere, and with considerable flexibility in my work hours. However, this is a two-edged sword as I can be in attractive locations with my family and can stay on top of my projects, but on the other hand, I almost never get away from my work. For a while, the trick was to take personal trips where I didn't bring my laptop. But today's smartphones are so becoming so capable that they let me stay all too connected.

This would all not be so bad except research suggests that the brain adapts to this sense of constant connectedness. The cover story of Time Magazine, January 29 this year, "The Brain - A User's Guide", made it clear that how we choose to live our lives, and what we mentally dwell on, actually shapes our brains. The article titled "How the Brain Rewires Itself" states "ý has come the realization that the adult brain retains impressive powers of 'neuroplasticity' - the ability to change its structure and function in response to experience." So the more connected we are, the more we adapt to that condition, and the less we can tolerate any delays in information. Otherwise, why would we all crave push e-mail that delivers our e-mail in seconds, versus much simpler polling approaches that might occur every fifteen minutes.

The New York Times on February 17 in a short piece called "Wireless Codependency", reported Robert Bornstein, a psychologist at Adelphi University, as saying "The superconnected may develop a dual-dependency. They're not only counting on other people too much, they're also hooked on the devices themselves, sometimes to the point where they feel utterly disconnected, isolated and detached without them." I guess we, the authors and the readers of this column probably all belong to this "superconnected" category.

My conclusion is that the more time we spend connected, especially if we allow a continual stream of asynchronous interruptions, the more fragmented our lives, and the less able we are to focus on the task at hand, and to enjoy the moments of our lives when we are not working.

The networks, devices and applications are here to stay. However, I think we all can make conscious choices about how most effectively to apply these technologies to our business processes, as well as to our own work and personal lives. For instance, I've turned off automatic delivery of e-mail to my inbox. Now I find that I can concentrate on work I'm doing, such as writing columns in intense and highly productive sessions. When I'm ready for a break, I manually click for new e-mail, and then respond to it as a diversion from the projects I'm working on. It baffles me how many people I see taking work calls on chair lifts when I go skiing. Are the calls so important that they warrant the disruption of all too few days of relaxation. I supposed the superconnected feel anxiety about any number of hours that go by where crises may be piling up in their inboxes.

So, I believe it's time to judiciously disconnect. This is harder than it seems, for I must confess that I've let myself go quite far down the connected path, partly because it's part of my job definition, so it's going to take quite a few disconnected hours to return to some level of normalcy.

EV-DO Rev A: Upping the Ante

Network Computing Mobile Observer, February 7, 2007

Both Sprint/Nextel and Verizon Wireless have just announced the availability of their Rev A versions of CDMA2000 EV-DO technology, which significantly boosts uplink speeds. This technology is the logical next step in the evolution of wide-area wireless.

On Jan. 30, Sprint announced it's CDMA2000 EV-DO Rev A (Code Division Multiple Access 2000 Evolution Data-Optimized Revision A) network as now available in south Florida; Portland, Oregon; and Puerto Rico, adding these areas to coverage that encompasses 95 million people in 21 other major markets. Sprint calls the service its "upgraded Mobile Broadband Network." Rev A is an upgrade to Rev 0, which became available in 2004, and ups the ante in wide-area wireless by presenting the highest performance claims yet for any available wide area wireless technology.

Specifically, Sprint/Nextel indicates a significant boost in typical uplink speeds that increases previous EV-DO Rev 0 speeds of 50 Kbps to 70 Kbps to a range with Rev A of 350 Kbps to 500 Kbps. Improved

uplink performance is no surprise as it has been on technology road maps, including ones I have published, for many years. What is a surprise to many in the industry, though, are greatly increased throughput speed claims on the downlink. Sprint/Nextel originally projected 450 Kbps to 800 Kbps for Rev A, but in last week's announcement the company indicated average transmissions between 600 Kbps and 1.4 Mbps. When a Network Computing colleague asked Sprint/Nextel about this, the response was that these faster speeds more accurately portray the range of speeds customers are likely to experience, based on Sprint/Nextel's actual performance tracking.

These values are quite a bit higher than what one would expect based on the increased spectral efficiency of Rev A over Rev 0, so if Sprint/Nextel can deliver these higher speeds, it may have more to do with improved transport technologies between the base station and the core network, rather than the radio technology.

Not to be outdone, on Feb 2, Verizon Wireless announced its Rev A network as being available in the greater Chicago area. I anticipate Verizon Wireless will upgrade the rest of its network this year. Its throughput claims, however, are much more modest, at 450 Kbps to 800 Kbps on the downlink and 300 Kbps to 400 Kbps on the uplink. Cingular Wireless, with its HSDPA (High Speed Downlink Packet Access) network,0 quotes 400 Kbps to 700 Kbps typical download speeds
with bursts to 1 Mbps or higher. HSDPA and EV-DO are both highly optimized CDMA technologies employing many of the same radio mechanisms, so at the end of the day, I would expect all these technologies to provide similar performance. Keep in mind that there are no rigorous definitions for "average user throughput," as many variables, including number of users in the cell sector (both voice and data), types of applications being used and average user signal quality,
are involved. So all these numbers are somewhat slippery.

It's important to note that enhancements for wireless technologies come bundled in major releases that necessitate network upgrades and generally new devices. In this respect, there has been a long-term leapfrog game going on between the GSM/UMTS (Global System for Mobile Communications/Universal Mobile Telecommunications System) and the CDMA2000 community. It began with GPRS (General Packet Radio Service) for GSM, which was top dog until CDMA2000
1xRTT (One Carrier Radio Transmission Technology) came along. Then Enhanced Data Rates for GSM Evolution (EDGE) jumped past 1xRTT. Then EV-DO Rev 0 leaped forward, followed last year by HSDPA. Now Rev A makes claim to being the fastest technology. However, HSDPA this year will be enhanced with a technology called HSUPA (High Speed Uplink Packet Access), and HSDPA
combined with HSUPA will be called simply High Speed Packet Access (HSPA). HSPA
should match Rev A in both downlink and uplink capabilities.

Enough about competition and claims. There are some other important aspects to Rev A worth understanding. One is that with a faster uplink, combined with QoS mechanisms in the network, it's feasible for carrier-provided VoIP. This is on the road map for CDMA2000 operators, but due to the complexity involved, I expect it's still a couple of years out. Rev A also reduces latency from

the typical 250-msec value for Rev 0 that I have measured, but I haven't seen any numbers published yet. This will help applications, especially chatty ones, as much as increased throughputs.

Rev A devices will include PC Card modems, USB modems and embedded Express card modules. Usage pricing has not changed. Sprint is charging $59.99 per month for a two-year contract, and a usage based plan of $39.99 for 40 MB, which is consistent with other operators. These prices still are a major obstacle for broader adoption, in my opinion, but more users at lower price points clogging up the network are apparently not in operator best interests. That's supposedly where WiMAX will come in, letting Sprint/Nextel leverage spectrum holdings much greater than it has for CDMA2000 to make wireless broadband available at lower price points.

Bottom line, I'm excited by the continual improvements in wide-area wireless technology, and Rev A is an essential step to a ubiquitous mobile broadband future.

IEEE 802.11n for the Enterprise—Not a Trivial Upgrade

Network Computing Mobile Observer, January 17, 2007

IEEE 802.11n will bring huge performance gains for Wi-Fi. Unless you've planned carefully, however, deployment in the enterprise will not be as easy as just upgrading or swapping out access points.

Six hundred megabits per second! That is the highest theoretical speed of IEEE 802.11n if a product implements all the mandatory, as well as optional, capabilities of the standard. The same core radio technology that enables these high rates also extends range, which results in increased AP (access point) coverage. Although speed and range are generally good things, there are ramifications, along with various other things to consider, with this standard. I bring this topic up now for two reasons. First, IEEE is planning to release the second draft (draft 2.0) of the IEEE 802.11n standard this quarter, with compliant products available shortly thereafter. Second, I'm teaching a course on Wi-Fi technology evolution and integration at Portland State University on January 26 (see http://www.rysavy.com/training.html for more details), so I've been busily creating and updating 802.11n slideware.

While this is a good time to start understanding migration to 802.11n, in most organizations it's probably best not to start deploying until the ratified version of the standard is available. And it's not expected until 2008. But you can expect tremendous promotion of draft 2.0 products this year, particularly because the Wi-Fi Alliance will be certifying products for interoperability. I personally find the concept of certifying against a draft standard a little strange, but the argument is that the draft standard should be reasonably close to the final standard, that it already incorporates a tremendous amount of scrutiny and review, and that the draft standard offers powerful capabilities. Thus, the certifications will help develop the market faster as well as offer customers benefits sooner. For instance, certified cards will interoperate with certified access points. So, certainly for smaller deployments like consumer and SOHO, the draft 2.0 products will make sense. Although it is hoped that

products will be either firmware upgradeable to the final standard or interoperable with final-standard products, there is no guarantee. In my view, deploying draft 2.0 in larger networks entails considerable risk, especially given the tortuous history of this standard.

Regardless of when you deploy, it's important to understand the speed and range capabilities so you can start planning now. I quoted 600 Mbps as the highest theoretical rate mostly to get your attention. This is indeed possible with 4 X 4 MIMO (Multiple Input Multiple Output with 4 transmitters and 4 receivers), a 40-MHz channel, a short clear view of the AP and some optional features implemented. But it is only barely theoretically possible. The actual mandatory requirement is 2 X 2 MIMO, which delivers a highest rate of 300 Mbps if all other features are implemented. So in the real world, as an educated guess, 2 X 2 MIMO throughput rates should turn out to be 100 Mbps to 200 Mbps, though actual rates across different types of environments remain to be seen. Still, we're talking a massive increase in speed compared to 802.11a or 802.11g, which max out at 25 Mbps of effective throughput.

You should think about how much speed you actually need. If you're surfing the Web, you'll likely be constrained by your Internet connection. But for sending video streams around homes, the bandwidth will be extremely useful. In work environments, higher throughput translates to more users per access point and better operation of services like VoIP, though these need a variety of additional enhancements such as QoS and fast roaming to be effective.

Where 802.11n speed impacts planning is with your Wi-Fi infrastructure. First, any decent 802.11n AP should have a Gigabit Ethernet connection. This may seem obvious, but some early products only supported a 100 Mbps connection. Next, if you're using a Wi-Fi controller architecture, can the controller support APs operating at much higher throughput rates? There may be backplane or processing limitations. Interestingly, some vendors--Trapeze Networks, for example--are moving some functions, such as cryptographic processing, back to the edge. This reverses the recent trend in the industry to centralize as much functionality as possible. At 802.11n speeds, it doesn't necessarily make sense for every packet to have to flow through the controller. So you'll need to see what precise migration plans your vendors may have for supporting 802.11n. While you're at it, ask them if access points can be upgraded to 802.11n by replacing radio modules or whether you'll need entirely new APs.

Like the previous IEEE 802.11a/b/g standards, 802.11n provides a tradeoff between speed and range. The highest speeds will only be available at short range. Although consumers, and eventually metro Wi-Fi networks, will love the improved range, greater range could cause interference between AP coverage areas that previously didn't interfere with each other. So you may need to reduce the power levels at which APs operate. This shouldn't be too complicated, but it does mean that just upgrading an access point by itself could cause problems.

Finally, there's the question of frequency planning. I generally recommend 802.11a for enterprises because of the larger number of channels and fewer interference sources. Fortunately, 802.11n can operate in both 2.4 GHz and 5 GHz, though it's up to vendors as to which frequencies their products support. If you plan to use the double wide radio channels of 40 MHz, realize that this makes frequency

planning more complicated--especially at 2.4 GHz, which supports three simultaneous normal width channels but only one 40 MHz channel. And while 40-MHz channels boost an individual's throughput rate, overall AP capacity does not increase.

Although I'm excited about 802.11n, I see most of the immediate benefits for consumers. Down the road, 802.11n will become the norm for larger deployments in organizations and municipal networks. But for the various reasons I've pointed out in this column, upgrading to 802.11n won't be trivial. So it makes sense to start planning now.

Google Embraces Mobile

Network Computing Mobile Observer, December 20, 2006

Until now, accessing the Internet with mobile devices has been an exercise in frustration. Sure, using the Internet to get to e-mail has worked fine, with a lot of different solutions available. But browsing the Net with pages formatted for large displays, or using various Internet-hosted applications, has not been practical. Luckily, things are changing.

In a series of moves this year, including last month's J2ME (Java 2 Micro Edition) client of GMail, Google has shown its commitment to the mobile space and what can be achieved by well-designed mobile applications. The J2ME client provides fast and convenient access to free GMail accounts; it also offers a much better user experience than using the Web-based approaches available from competing services such as Yahoo.

Another application Google has adapted for mobile use is Google Maps. I've already found Google Maps on a desktop to be a super application, especially with its satellite photo integration and the ability to provide directions. Mobile client code is currently available as either a Java application or a native application for the Treo. I've been using the mobile maps applications for a couple of months, and find it an extremely useful mobile application. You can drag the stylus to move the map, and easily zoom in and out. In an unprecedented acknowledgement to the needs of wireless computing, the application shows in real time how much data is being consumed. This feedback not only gives you a sense of why the screen may be updating at the rate it is, but for anybody on a usage-based plan, it tells you how quickly you are consuming your bucket of bits. It's easy to go through hundreds of kilobytes per minute, and any prolonged use goes into the megabytes. Anybody using the application on a regular basis may want to consider an unlimited data plan, which fortunately is less expensive for phone than laptop usage. The application also provides real-time traffic information. Once the expected integration with mobile handset location information occurs, the application will be even more useful.

Google also offers a customizable mobile homepage featuring a search box and items such as weather and personalized news. I'm particularly impressed by the mobile search capability, in which you can specify searches across the Web at large or on just mobile sites. For general Web searches, if you click

on a search result, Google then works as a proxy for the response, rendering it in a format more suitable for a mobile screen.

Finally, Google allows you to use SMS messages to obtain specific types of information, including local TV listings, weather, sports, movies, driving directions, area codes, calculations and more.

Clearly, Google has seen the potential of mobile computing and intends to be a driving force in delivering personalized, localized and efficient content and applications. The general quality of results is far higher than most of the operator-hosted portals, which I've generally given up on. To give appropriate credit, other Internet players are also targeting the mobile environment, among them AOL, MSN and Yahoo. But so far, I've been most impressed by the scope and execution of Google's efforts.

What does this mean for an average enterprise application of wireless? Well, it's hard to predict all the consequences. In general, however, the broad applicability of mobile and wireless technology for both business and consumer markets has driven down costs and made the technology universally available and affordable. Well-designed consumer-oriented applications will no doubt drive business-oriented applications, and some of the applications I've described here--such as mobile-optimized searches and maps--can be used in work scenarios today.

It's good to see a company like Google take such sensible steps to target the mobile market opportunity. Granted, these are just initial and somewhat exploratory steps, much like Google's metro Wi-Fi network in Mountain View, Calif. Also, the Mobile Web still only offers a small fraction of the richness of the Web at large. However, they point to a future of diverse and compelling mobile applications. With companies like Google taking the opportunity seriously, that future will arrive just that much faster.

Using VPNs With Wireless Networks

Network Computing Mobile Observer, November 29, 2006

I've been involved with multiple projects to assess the use of VPN (virtual private network) technologies over wireless networks. The most recent is pulling content together for a Webcast I'll be doing on December 5 for Cingular Wireless titled "Optimal Use of VPNs Over Cellular Networks." The good news is that an increasing number of effective options result in improved performance, reliability and control. The bad news is that all the options and tradeoffs are complicated.

Two items have driven progress. First, the networks themselves have become much faster, with current 3G networks delivering average speeds of over 500 kbps, and as of the end of 2006, they are fairly widely available. This development has helped networking applications in general, and VPNs specifically, because most VPNs were not designed for wireless and impose tunneling overhead in their additional packet headers. Second, a number of vendors have developed VPNs specifically for mobile operation, and these are now becoming extremely sophisticated with features such as traffic shaping.

Many organizations are already using VPNs for remote access, replacing their dial-up remote-access servers with systems that allow users to simply connect to the Internet from anywhere and then engage in secure (encrypted, authenticated, tamper resistant) sessions. The advantage is that both ends of the connection simply need an Internet connection. Remote users can connect via dial-up to their ISP, DSL, cable modems, Wi-Fi and, increasingly, wide-area wireless such as EV-DO or HSDPA.

There are a number of reasons you would want to use a remote-access VPN with a wireless network. First, you can't depend on the provider encrypting the radio link. Most public Wi-Fi networks operate in the clear. In the cellular world, many networks use encryption, but not all. And even for the networks that use encryption, it is usually to a node in the infrastructure beyond which data passes in the clear. Granted, this may be over a private operator network. But for sensitive data, this might still make you nervous. Most important of all, wireless network connections, be they cellular or public Wi-Fi, generally use the Internet to traverse from the operator network to your organization. By using a remote-access VPN, you can secure the communications on an end-to-end basis, you are not dependent on any of the security features of the underlying networks, and you can deploy a consistent security solution regardless of the access network.

But the question is what kind of VPN to use. There are three main categories, including IPsecIPsec VPNs, the ever more popular SSL VPNs and mobile VPNs. We'll quickly look at all three. IPsec VPNs are the workhorses in the industry, and many organizations have deployed them for remote access. They work fine over wireless connections, but they do add protocol overhead, with some 50 bytes per packet. For larger packets, as in a file download, this may not be that noticeable, but chatty applications may operate perceptibly slower. The performance penalty may range anywhere from 5 percent to 30 percent, depending on the type of application. This is less of a factor with 3G networks than 2G networks. Also, VPN sessions are vulnerable to connection loss, which can force users to restart their VPN as well as the applications they were running. In addition, you may need to configure the VPN for NAT traversal by enabling UDP encapsulation. Bottom line: IPsec VPNs work best with stable and fast connections.

An increasingly popular option for remote access is to use an SSL VPN, which leverages the Secure Sockets Layer found in most browsers. This allows clientless operation, though it restricts operation to Web-based and file-access applications. However, since many handheld devices have browsers with SSL, this type of VPN makes it relatively easy to support a wide range of mobile devices. Companies like Aventail have also created mobile versions of their SSL VPN products, where the security gateway shows a portal page formatted for small-screen devices and allows users to limit mobile device access to applications and content that make sense for the particular device. With the addition of client code, SSL VPNs can support a wide range of applications. Just make sure client code is available for the device of interest.

Finally, there are mobile VPNs, from companies such as NetMotion Wireless, designed from the ground up to handle the complications of mobile networking. These have become extremely sophisticated and now support a wide range of features, including protocol optimization with compression, session

maintenance when you temporarily lose a connection or suspend a device, roaming between different network types such as 3G and Wi-Fi, and new features like traffic shaping that let you prioritize traffic--giving VoIP higher priority over other tasks, for example, and even blocking some applications (such as Windows Automatic Update) when operating on slower connections. NetMotion Wireless announced this capability for its product earlier this month. Almost any wireless application involving frequent use and use while mobile will benefit from a mobile VPN. The tough part is if you have already standardized on an IPsec or SSL VPN for remote access, which means you might need to maintain two separate remote-access solutions.

Between improving VPN choices and faster wireless networks, using VPNs for wireless remote access has never worked better. Just make sure you use the right one for your situation.

Ditching the Laptop

Network Computing Mobile Observer, November 8, 2006

Symbian executives at the Symbian Smartphone Show in London last month spoke about the increasing role of smartphones, even suggesting that with the advances in smartphones, people may wonder soon why they need PCs at all. They pointed out that in developing countries where wireless penetration is far outpacing wireline penetration, a phone is a more logical connectivity appliance than a PC. The execs noted that in India the PC market is growing at 5 million units a year, while mobile phones are growing at that rate each month.

I think there is indeed a computing-paradigm shift happening, though I think we are at the earliest stages. With advances in computing power and storage, as well as high-speed wireless networks, it should be possible by the end of this decade to have a small handheld device that can run most of the applications you need, that can store much of the entertainment you care about--not just songs but full-length movies--and that has constant broadband connectivity. The only thing missing, and it's a big thing, is the user interface. Maybe by then we'll have retinal imaging and reliable voice input, though I'm skeptical. More likely we'll be able to easily dock our phones, perhaps using UWB, to conventional keyboards and displays in our homes, offices and maybe even hotels.

Certainly, we have hints of all these possibilities with current devices, applications and other innovations. Here are some of the developments that make it possible for many people, including me, to travel more often without a laptop on shorter trips. And as somebody who has taken a laptop everywhere for the last 15 years, this is a welcome change.

First is general ease of use. RIM recently introduced the BlackBerry Pearl, one of the easiest BlackBerries to use ever, especially with the new trackball. After using the device for some time, I have to admit that the trackball is the fastest and most effective navigation tool I've ever used on a phone. Traditional controls such as touchscreens and five-way navigation pads are also effective. The industry is still trying to innovate its way to efficient keyboards for small devices, but is doing reasonably well given the

extremely limited real estate. However, let's face it, typing speeds on a phone are always going to be much slower than on a full-size keyboard. For data retrieval and e-mail monitoring with occasional short replies, though, I find either the fully micro-QWERTY or the RIM SureType keyboards just fine.

Another huge area of innovation is in applications. There is no shortage of effective wireless e-mail solutions. But not as well recognized necessarily is the improving ability to deal with documents. I've spent some time working with the just-released version 9 of Documents to Go from Dataviz, which allows users to work with Microsoft Office documents on Palm OS or Symbian devices. I have found the product flawless in my fairly limited use, allowing me to both view and edit documents. It can also view PDF documents. RIM's file viewers provide similar functionality; although they do not allow editing, they more than making up for it by allowing selective viewing of what might otherwise be files too large to download over a wireless connection. And of course, Windows Mobile has long touted its Office integration. Other types of applications are also making their way to mobile devices. Google, for example, just completed an excellent adaptation of Google Maps.

Then there is mobile-formatted content. I've been pleased by how many publications now offer mobile-formatted content, including Time Magazine, Yahoo, The New York Times, Google, The Washington Post, BBC, The Wall Street Journal and others. Now I have something to keep me occupied no matter how long the line at the post office or the bank. Google's auto-mobile formatted content in response to searches is also very handy.

Storage is another big enabler. Last month I bought a 2-gigabyte Secure Digital card for just $40. Now I can have dozens of music albums with me, not to mention all the documents I would ever need. 4-GB cards are available, with no limits on eventual size that I'm aware of.

Finally, all work and no play makes a phone a dull toy. There are increasing music options for downloading over the air or transferring from a PC. With thousands of radio stations streaming around the world on the Web, it's quite easy to capture shows and transfer them to your phone for the ultimate radio place-shifting time-shifting experience. The inevitable direction here is for our phone to become fully capable digital video recorders, either capturing mobile TV content or receiving content from our home DVRs.

It doesn't take much effort to extrapolate from some of the examples I've given of what's already available today to realize the huge effect that the mobile phone is going to have on the entire computer industry. And, ultimately, it just may be as disruptive as the PC was to mainframes.

Evolved EDGE and the Future of TDMA

Network Computing Mobile Observer, October 18, 2006

Late last month I attended an analyst meeting in Boston hosted by 3G Americas, an organization that promotes GSM/UMTS-based technology in the Americas. One area of emphasis at the meeting was the likely evolution of 3GPP wireless technologies, including continued improvements in UMTS such as High Speed Downlink Packet Access and High Speed Uplink Packet Access (with the combination referred to simply as High Speed Packet Access or HSPA), enhancements to HSPA called HSPA+ and the eventual move by the end of the decade to 3GPP Long Term Evolution--an OFDMA (Orthogonal Frequency Division Multiple Access) technology that shares a number of attributes with Mobile WiMAX. There are no surprises in this area, but information about the continued evolution of technologies based on TDMA (Time Division Multiple Access), including Evolved EDGE (Enhanced Data Rates for GSM Evolution), was new to many attendees. Evolved EDGE is a technology I studied fairly closely as part of a white paper I developed for 3G Americas. Published last month, "Mobile Broadband: EDGE, HSPA and LTE" was distributed at the meeting and is available at http://www.3gamericas.com.

GSM with EDGE today is by far the most broadly deployed and used wireless voice and data service. Globally, GSM has over two billion subscribers, gaining a whopping 500 million of these during the 12 months from the second quarter 2005 to the second quarter 2006. As of September 15, 2006, 160 operators in 91 countries are offering commercial EDGE services. It's likely that most GPRS (General Packet Radio Service) networks eventually will be upgraded to EDGE. Some operators were originally skeptical of EDGE, with UMTS coming online. However, there is an emerging realization that UMTS--because it requires a sizeable investment at every base station-- won't be deployed everywhere for quite some time. Additionally, providers don't want to drop UMTS/HSDPA users all the way to GPRS when they roam out of UMTS coverage. As described by a wireless architect I know, it's kind of like a car that goes only 100 mph or 10 mph. EDGE provides better service continuity between the two services. Hence, even many UMTS operators that were never going to deploy EDGE are now doing so.

Evolved EDGE enters the picture because continued radio technology improvements now make it possible to take EDGE to a whole new performance level. One such enhancement is duplexer technology, which allows simultaneous send/receive. Current devices cannot send and receive simultaneously, thus limiting the number of time slots an EDGE device can use in the radio channel. Another enhancement is mobile receive diversity, where two radio receive chains operate in parallel on the device. Outputs are combined to extract a better signal, thus facilitating lower coding overhead and higher order modulation. Higher order 16-QAM modulation is yet another enhancement, as is simultaneously receiving on two radio channels. Add all these up and you get Evolved EDGE in 3GPP Release 7, which will have peak throughput rates of 1.3 Mbps with achievable user data rates of over 1 Mbps--previously the domain of technologies such as HSDPA. Of course, peak rates with new versions of HSDPA devices are now projected as high as 7.2 Mbps and, eventually, with HSPA+ as high as 28 Mbps. That's all well and good, but a super- charged version of EDGE actually could be very attractive for a number of operators.

First, there are quite a few GSM operators globally that have not yet committed to 3G deployment. These operators are doing just fine selling voice and lower-speed data services such as wireless e-mail. With the whole business model for mobile broadband somewhat up in the air, Evolved EDGE could be a

much lower-cost way of beefing up data service while waiting for the mobile broadband business to mature. Even for operators offering 3G service, Evolved EDGE could allow 3G deployments to be focused in higher-density areas while a lower-cost alternative for higher-speed data service is offered in rural areas. All this will be especially attractive because some Evolved EDGE features will be a relatively straight-forward software upgrade to the network. In addition, Evolved EDGE devices are not likely to cost much more than EDGE devices. For price sensitive market segments, a single radio solution based on GSM/EDGE will be less expensive than a UMTS device that must also support GSM/EDGE, thus requiring two radios. The one radio solution brings power consumption benefits, too.

So, even as most attention focuses on 3G, evolved 3G and Mobile WiMAX technologies, TDMA-based GSM--along with its data service EDGE--will constitute the majority of subscriber usage through at least the end of the decade. Efforts like Evolved EDGE also show that it does not so much matter what the underlying radio technology is, but what service and benefits it can offer users.

Mobile Computing Policy and De-Perimeterization

Network Computing Mobile Observer, September 27, 2006
By Peter Rysavy

On September 20, I chaired a meeting of the Portable Computer and Communications Association (pcca.org) on the topic of "Mobile Computing Policy and Network Access Control." The PCCA meets quarterly to address developments in wireless and mobile computing, and meetings have good representation from operators, device vendors, computer vendors and wireless middleware providers. This meeting, hosted by NetMotion Wireless, proved quite illuminating, making me realize that mobile computing is simultaneously maturing and becoming an evermore complex field, with new aspects to consider. Policy management is one aspect. De-perimeterization (not an English word, but used at the meeting and a descriptive nevertheless) is another.

Policy refers to rules on how computers are allowed to behave in different scenarios. For example, Microsoft Windows Vista Group Policy allows IT managers to specify items such as which SSIDs (Service Set Identifiers) a laptop computer can associate with, and whether to prevent a connection if WPA (Wireless Protected Access) isn't available. Mobile computing policy is a capability in products such as NetMotion Wireless' Mobility, where IT managers can centrally specify rules such as what networks a mobile system can connect with and, more interestingly, what applications are allowed access to what network. This way, bandwidth-hogging applications (e.g. music downloads) might only run over Wi-Fi or Ethernet, but not over a 3G connection, which could be especially handy with a usage-based pricing plan.

Since the NetMotion Wireless product sits above the stack, it actually knows which applications are transmitting information and can go one step further and do traffic prioritization within the secure tunnel that the mobile VPN has created between the client software and the mobility server.

For instance, the software can give higher priority to VoIP traffic than Web browsing, while giving Web browsing higher priority than Windows Update operations. The result is a form of QoS in networks that currently have no QoS capability, and the results are impressive. A demo showed intelligible voice over a 3G connection with other applications running, compared to completely garbled voice.

What's nice about a well-implemented policy management system is that rules can be created centrally and pushed out to devices. At the meeting, Aventail provided another example of policy in describing its Mobile SSL VPN product, which enforces endpoint control. Here the central SSL VPN gateway can interrogate the device to make sure it has the right constitution before permitting it to access enterprise resources. For example, with Windows Mobile 5 devices, the system can check for specific applications, directory names, file names, Windows registry entries, Windows version and device certificates. You can expect more and more mobile middleware solutions to support policy management, as well as increasing support from within the OS itself.

As for de-perimeterization, Aventail refers to this as an inverted network and promotes this approach with its SSL VPN solution. The concept is simple. In the past, companies designed their networks to have hard perimeters, with little protection once you were inside the network. The concept of an inverted network is to trust no node and to require every point in the network to have authenticated/authorized/encrypted access to enterprise resources. This adds some overhead to the network, but it offers quite a few advantages. It greatly facilitates guest network access, temporary workers using your network and workers connecting to your network by the network connection of the moment, whether 3G, Wi-Fi hotspot, home network or Internet kiosk. In other words, it enables a flexible networking topology sensitive to today's increasingly mobile paradigm.

At the same meeting, Boeing described its implementation of an inverted mobile network architecture based on the Open Group Secure Mobile Architecture design. The core networking protocol used within SMA is Host Identity Protocol, a variant of IPSec that addresses mobility by relying on names rather than IP addresses to identify endpoints. This enables roaming across subnets and is an elegant alternative to using the Mobile IP protocol. HIP is currently an IETF draft RFC. The Boeing mobile architecture also implements policy in the form of location. Mobilized equipment connected over Wi-Fi can only communicate when it's in an authorized location. I found the Boeing mobile network architecture quite impressive, and it's no surprise given Boeing's active involvement in key underlying standards and wireless technology evaluation. My first discussions with its wireless architects date back to 1993.

Policy and de-perimeterization are just two instances illustrating how mobile and wireless technology implementation is much more than just a discussion of wireless links. It's easy to get fixated on the wireless link, as wireless technology is so sophisticated and so interesting. But ultimately, that wireless link just replaces a wire, and a complete, managed and secure system for mobile workers requires many ancillary components.

Trials and Tribulations in Assessing Wireless Network Performance

Network Computing Mobile Observer, Sep 6, 2006

I've been involved in several projects assessing wireless network performance over the last year, and I've gained an appreciation for the associated complexities. Testing has included both Wi-Fi hotspot networks and 3G networks. 3G network testing involved both CDMA 2000 EV-DO (Evolution Data Optimized), as offered by Sprint and Verizon, and HSDPA (High Speed Downlink Packet Access), as offered by Cingular. My most recent effort involved reviewing and advising on the testing done for the August 31 Network Computing cover story I wrote, "Taking Advantage of Wide-Area Wireless--It's Time to Decide"

One of the first factors you must consider is that in the past, the wireless network was usually the bottleneck. However, today's wireless networks are becoming so fast that other parts of the network may be the constraining factor, such as the backhaul between base station and operator core network, an item I wrote about last time when commenting on Sprint Nextel's WiMAX choice. Another bottleneck could be the speed of the Internet connection of the server you are accessing while mobile. Even if the speed of that connection is faster than the wireless connection, there may be other traffic over that connection limiting the amount of available bandwidth. Or there may be security gateways that can't keep up with the remote access network. These back-end connection issues actually forced us to redo quite a few of our tests for last month's Network Computing story, which initially showed HSDPA throughputs to be quite low (approximately 300 kbps) but then considerably higher (700 kbps to 1 Mbps) once we eliminated the server-side bottlenecks.

In a comparison of Wi-Fi hotspots vs. EV-DO--published in a free white paper available on my Web site--we took no chances with back-end limitations and tested against a dedicated server that we collocated at an Internap collocation site with a 100 Mbps connection. For throughput, I prefer doing FTP (File Transfer Protocol) tests using test files of 1 megabyte to 3 megabytes. It's also possible to use public Web sites that provide throughput measurements, but I don't consider these as dependable or as repeatable as testing against a known and controlled test server. Even with file transfer, however, there are complications. One of our tests was to load the network with four simultaneous downloads. With EV-DO we found it difficult to get the four downloads to occur at the same time; the latency of EV-DO caused problems when attempting to have client computers try to open the FTP data port at the same time. So we switched to using HTTP Get operations, which worked much better because of the simpler protocol involved and allowed us to demonstrate how just four simultaneous downloads on an EV-DO cell sector can affect throughput. The average throughput we measured for a single computer--430 kbps--went down to 248 kbps with four systems using the network.

File download tests measure one aspect of a wireless networks capability, but they don't capture aspects such as latency. One good test here is to load Web pages that are representative of Internet Web sites, for example, complex pages such as www.cnn.com. The challenge here is that public sites vary their content moment by moment and have varying usage loads. In our EV-DO vs. Wi-Fi hotspot

testing, we solved this problem by building our own Web page on our test server with the same complexity as cnn.com, which has about a 100 graphical objects. With Web downloads, it's important to clear the caches before each download, otherwise only a portion of the page may be downloaded in subsequent tests. In our results, the hotspot network averaged 6 seconds for downloading the page versus EV-DO, which averaged 25 seconds, demonstrating how latency can affect some applications more than throughput. Another good test is to send and receive a certain number of e-mail messages, as this also is representative of real-world operation and captures the effects of both throughput and latency.

Another aspect of testing networks, especially cellular networks, is that voice and data loading can affect performance, and performance varies depending on signal strength and interference. Hence, it is not enough to measure in just one location. In fact, you need to measure in at least eight locations to obtain representative results. More locations would be even better, but costs really start running up. To complicate matters further, operator networks may use one vendor's radio-access network hardware in one market and another vendor's equipment in another market, with possibly different performance. I have seen significantly different throughput rates in different cities with some operator networks.

Measuring time is also a problem. For ad hoc testing you can use your watch, but for operations taking seconds, this method is not very accurate. In our testing, we built test scripts that would invoke all test operations and automatically capture the amount of time involved.

Measuring latency is not trivial either, as we learned in our recent HSDPA testing. In the past, our tests would use 11 successive pings, discarding the first one because that one can take extra time to bring up the data traffic channel on cellular networks. With HSDPA, however, we saw much higher latency results than we expected--300 msec to 400 msec. It was only after conferring with the operator that we learned that it takes a certain amount of traffic to invoke the high-speed data channel. Our ping packets traveled on a lower speed channel and were not representative of what an application would actually obtain. We needed to increase the ping test rate to twice a second. The DOS command-line ping command in Microsoft Windows only does pings once a second, but you can use a utility such as PingPlotter to increase the ping rate. Once we did this, latency results went way down, to as low as 130 msec when testing against google.com.

Bottom line: It's easy to do some quick and dirty tests to characterize a wireless network's performance. However, to comprehensively characterize these networks is a major undertaking.

Sprint Nextel, WiMAX and the Mobile Broadband Conundrum

Network Computing Mobile Observer, Aug 16, 2006

Sprint Nextel announced on August 8 its choice of WiMAX technology for its mobile broadband network that will operate in its 2.5 GHz Broadband Radio Service (BRS) spectrum. Every other wireless pundit is going to have something to say about this, and I have so much to say that the constraints of this column

will be a challenge. First, I'm going to make some general comments about the company's technology choice and early claims. Then I'm going to zoom in on the crux of the challenge that Sprint Nextel will have to address: managing capacity and performance expectations while finding a way of charging for service that enables a successful business model. This will be a conundrum, which according to one definition I just read means "a paradoxical, insoluble or difficult problem."

To begin with, it's clear that Sprint Nextel's choice of Mobile WiMAX (based on the IEEE 802.16e-2005 standard) is a huge win for Intel--especially in conjunction with Clearwire, which also recently chose WiMAX. The technology has now achieved some credibility. Mobile WiMAX embodies many radio innovations and, on paper at least, has impressive capabilities. I do take exception, however, to labeling WiMAX a 4G technology. First, there is no official 4G standards work. But people have been researching 4G approaches for quite a few years, and commonly accepted requirements include 1 Gbps peak speed and 100 Mbps average speeds--50 times faster than what Sprint Nextel will be able to provide. However, there is some karmic balance in all of this, because Qualcomm, a WiMAX opponent, has been promoting CDMA2000 1xRTT for many years as 3G. 1xRTT has a peak network speed of 153 kbps. If 1xRTT is 3G, then that would definitely make Mobile WiMAX 4G. But the fact is, Mobile WiMAX will only yield a modest improvement over 3G systems of today, and the aggressive 3G roadmaps I described in my column three weeks ago will allow these systems to largely match Mobile WiMAX performance.

Now, let's get to the conundrum of capacity/performance/pricing. The first aspect is one of capacity. Regardless of how fast the radio technology is, most cell sites today in the United States are backhaul constrained by T1 circuits. A network architect commented to me last week, "Mobile WiMAX. That's just another fast airlink. When will people realize that the real issue is the wireline network?" Think about it. If a Mobile WiMAX system uses 10 MHz radio channels, and assuming 7.5 MHz is allocated to the downlink (this is a time division duplex system) and using a spectral efficiency value of 1 bps/Hz/sector (an optimistic value in itself for a loaded network with lots of interference), this is 7.5 Mbps of throughput per sector. With three sectors in a cell, this would require a whopping 15 T1 circuits. Clearly, this is not going to happen any time soon. Eventually, yes. 2008? No. This takes me to the next issue, which is performance. Sprint Nextel has indicated 2 Mbps to 4 Mbps average throughput. I view this as a very aggressive claim. First, the company will need all the backhaul capacity I described to support these rates. But even if Sprint Nextel had it, how many active users in a cell sector (one mile is typical cell site spacing) could the company support? Simply divide 7.5 Mbps by 2 Mbps to 4 Mbps, and you end up with a small number. Granted, not everybody is downloading at the same time, so the network can be oversubscribed, but only by so much before things slow down below the promised rates.

These are the same types of issues that cellular operators face. But what's different with 3G is that much of the data use is with bandwidth-efficient applications such as SMS and RIM BlackBerries. Business data users tend not to consume under 100 megabytes per month on average, so capacity has not been as much of an issue. But Sprint Nextel is emphasizing the consumer front, which is multimedia driven. I've seen dimensioning exercises for Mobile WiMAX with figures of 1 gigabyte per subscriber per month. Compared to 3G, this is lots of capacity--in fact, about 10 times the amount of a 1,000-minute voice

user. But the problem is that 1 gigabyte is less than one DVD-quality movie download, even using MPEG4 compression.

And that's the conundrum. 3G networks can easily support voice and pictures, so phone service, e-mail and Web browsing are no problem. But throw in multimedia, and you rapidly start using up all your capacity. 3G operators have dealt with this in their unlimited-use pricing plans by disallowing these bandwidth-hogging applications in the fine print of their service agreements. But those are the very applications that consumers want. Usage-based pricing is one solution, but the market has spoken against this approach. Sprint Nextel is going to have to solve the problem of how to make money while providing sufficient bandwidth to users in a network that has relatively finite capacity. Yes, the company has some 100 MHz average capacity per market, so that 10 MHz radio channel I described could eventually become 10 radio channels. But that doesn't change the fact that there will be a certain dollars per gigabyte that it costs to deliver capacity as well as a certain number of gigabytes that users will expect to access and only so much they will be willing to pay. Bottom line: The technology choice was easy compared to what Sprint Nextel will face in the business equation. Success could allow the company to establish a leadership position in mobile broadband in the United States, but it will be a high-risk undertaking.

Wide-Area Wireless--The Next Five Years

Network Computing Mobile Observer, Jul 26, 2006

Recently, I concluded the research phase of a large project to assess the future direction of all the major wide-area wireless technologies, including 3G and WiMAX. There is so much going on that it has left my head spinning. But at the same time I'm quite excited, because the future of mobile data just keeps getting brighter. As an IT manager, you won't be able to immediately take advantage of many of these enhanced technologies, but it may be helpful to know what will be available in what timeframe. It's also interesting to see how the 3G vs. WiMAX battle is shaping up.

3G today consists of the GSM/UMTS technology family and the CDMA2000 family of technologies. The version of UMTS (Universal Mobile Telecommunications System) that Cingular and other operators around the world are deploying is based on 3GPP (Third Generation Partnership Project) specifications release 5. This release includes a service called HSDPA (High Speed Downlink Packet Access). HSDPA has theoretical rates to 3.6 Mbps with today's devices, increasing to 7.2 Mbps in 2007. Real-world limitations, however, limit the peak rates you will experience to just over 1 Mbps. With multiple users active in a cell, speeds may drop to the 500 kbps range. Cingular is quoting typical rates of 400 kbps to 700 kbps. Interesting, it's only in the United States that operators quote "typical" rates. These are actually very slippery numbers, as they are based on backhaul architecture and capacity, cell site spacing, voice load and data load. Beyond HSDPA, 3GPP Release 6 specifications include a new service called HSUPA (High Speed Uplink Packet Access), which boosts peak uplink speeds to just over 1 Mbps, with theoretical peak rates of 5.76 Mbps. HSUPA will be available next year. People are calling the combination of HSDPA and HSUPA simply HSPA (High Speed Packet Access).

However, with UMTS/HSPA, it doesn't end there. A series of improvements are planned for HSPA in a combination of efforts. The first will come in the context of Release 7 functionality, and the next will come in an effort called HSPA Evolution or "HSPA+." Through a combination of radio techniques including channel equalization, mobile receive diversity and MIMO (Multiple Input Multiple Output), HSPA+ takes CDMA close to maximum theoretical efficiency. Peak rates could exceed 25 Mbps in a 5-MHz downlink/5-MHz uplink spectrum allocation and users could experience two to four times the rates they can with HSDPA today. This technology, which could be ready by 2008, could match mobile WiMAX in the same amount of spectrum.

Finally, to round out their offensive strategy, 3GPP is also busily defining an OFDM-based system called 3GPP Long Term Evolution, slated for initial deployment in the 2009 timeframe, with peak rates of 100 Mbps in a 20-MHz radio channel. The combination of HSPA+ and LTE will competitively squeeze WiMAX. Mobile WiMAX on paper exceeds HSPA performance, but it won't necessarily exceed HSPA+ by very much. Meanwhile, LTE is designed to be more efficient than Mobile WiMAX. With respect to timing, initial Mobile WiMAX networks could start appearing by the end of 2007 but are more likely in the 2007 timeframe, assuming operators embrace the technology and find spectrum to deploy it in.

If WiMAX doesn't have enough competition, it also has to watch the aggressive development moves of the CDMA2000 camp, as represented by 3GPP2 (Third Generation Partnership Project 2). The first improvement to today's EVDO (Evolution Data Optimized) Rev 0 technology will appear next year with Rev A, which boosts peak downlink throughputs from today's 2.4 Mbps to 3.1 Mbps, and boosts uplink throughputs from today's 153 kbps to 1.8 Mbps. Like HSPA, peak rates that users experience in either direction will likely be around 1 Mbps, with lower average rates. For Rev A, Sprint has indicated 450 kbps to 800 kbps average for downloads and 300 kbps to 400 kbps average for uploads. Beyond this, and available by 2007 for deployment though no operators have yet committed to this, EVDO Rev B can combine up to 15 1.25-MHz radio channels in 20 MHz of spectrum for peak rates of 73.5 Mbps. More likely, an operator would combine three channels in 5 MHz of spectrum for peak theoretical rates of 14.7 Mbps on the downlink, matching HSPA peak throughputs. 3GPP2 is also working on Rev C, which is likely to be an OFDM-based approach. One avenue here is the possible convergence of Rev C with the IEEE mobile broadband standard, IEEE 802.20. 802.20 is currently on hold because of allegations of impropriety in the standards process, but this won't necessarily affect the long-term outcome.

WiMAX proponents will have to try and prove that their technology is sufficiently better, and sufficiently less expensive, than alternatives to gain acceptance. Large companies like Intel, Motorola and Nortel are all betting heavily that they can do this. My view is that the massive investment involved in R&D and deployment--and the competitive pressures from other approaches like metro Wi-Fi--makes it likely that only two wide-area technologies can survive long term. I don't see GSM/UMTS/HSPA/LTE going away because of its dominant global market share today, meaning that five to 10 years from now, people may be trying to remember what the acronyms EVDO or WiMAX stood for.

Anatomy of a Well-Designed Wireless Application

Network Computing Mobile Observer, Jul 5, 2006

Over the last month I've been evaluating some new wireless e-mail products, and it has made me think about what it means for products to be optimized for wireless. After a decade of working with wireless applications, I thought I knew all the tricks, but it turns out that there is no shortage of ongoing innovation when it comes to this medium.

I've stated in previous columns that off-the-shelf IP-based network applications generally work quite well over today's new 3G networks, such as EVDO and HSDPA, especially for lighter use. But for workers who use the applications hour by hour for work-related functions, it is very important to evaluate how reliable the application is and how it deals with inevitable wireless effects such as weak signals (which reduces throughput and increases latency), occasional disconnects and possible changes in IP address when reconnects do occur. Today's mobile middleware addresses these effects quite well. The purpose of a well-designed wireless application, however, is to address these issues directly and then some. It's instructive to look at examples.

The two e-mail products I've been working with are RIM's new 8700 and a Palm OS Treo application called ChatterEmail. Both address the smartphone form factor and both do wireless e-mail extremely well. Whereas RIM's solution is ideal for organizations with larger numbers of users because of its extensive management features, ChatterEmail is essentially an end-user product that emphasizes Internet mailbox access using either POP 3 (Post Office Protocol 3) or IMAP (Internet Mail Access Protocol). As it turns out, IMAP provides much richer functions than POP3. ChatterEmail is also in the process of making an Exchange version available.

The first obvious feature is the ability to push information, meaning that new e-mail (or other enterprise information in the case of RIM) appears automatically on the new device without any action taken by the user. RIM does this via its Blackberry Enterprise or operator-hosted server and its network operations center. ChatterEmail does this by using the IDLE command that is part of IMAP, accomplishing push without the need for an intervening server. Push is great for an always-on smartphone. For laptop applications, it might or might not make sense based on how users interact with the application.

The next feature is less obvious, namely having the application work diligently in the background to obtain information, even in the presence of network anomalies, without bothering the user with the details. RIM does this brilliantly, since it has control of the end-to-end (device to NOC) connection. From the user perspective, the e-mail just shows up, no matter how many retries may have occurred. ChatterEmail only has control of one end, resulting in a reduced ability to address networking problems. It compensates for this cleverly by doing multiple retries, even giving the user options in the settings about how aggressive it should be in attempting to connect or whether to even try in the presence of a poor signal. Before you think that maximum aggressiveness is a good thing, realize that this can result in

high power consumption. ChatterEmail provides a quick command to reset network connections, which often resolves issues if e-mail is not getting through. It also provides a status in the corner of the screen of the connection steps it is going through so you can more easily diagnose issues in the case of failures.

Another area is power consumption. The RIM solution is extremely efficient in its use of the network, which helps maximize battery life. RIM can do this thanks again to control of the end-to-end connection. ChatterEmail doesn't have this luxury, as it's using an existing networking protocol that wasn't designed for power efficiency. So in its online state where e-mail appears immediately, battery consumption is relatively high, though good enough to get through a full heavy e-mail work day. But again, it compensates by having a more battery efficient polling mode, with user programmable intervals. This takes away from immediacy, but for users in meetings or in the evening, this may be just fine. A nice little feature here is that the program indicates in the
corner when it will next retrieve e-mail.

The next important area is how to deal with large amounts of information that you don't necessarily need on a mobile system, laptop or smartphone. With e-mail, this generally applies to large attachments. RIM handles this nicely by allowing you to view attachments piece by piece, so you may only download a few kilobytes to see what is in an attachment, versus megabytes of the original presentation file. Again, this requires control of both ends, with the Blackberry server intelligently dishing out portions of the attachment. ChatterEmail can't do this because it doesn't have the server component, but it cleverly provides other options for dealing with large e-mails. First, it lets you view the message without downloading the attachment; then it allows you to move the message into an IMAP folder on your server without downloading it. This way, you can keep your mailbox organized until you can deal with large messages and attachments using a full-screen computer. Both RIM and ChatterEmail provide options for deleting messages on both the device and the server.

I've only touched on the main features in this column. The bottom line is that applications tailored specifically for the wireless environment can deliver an extremely effective usage experience. Keep in mind, however, that what makes sense on a smartphone -- because of limited battery power and the device's always-on status -- may not be the optimal behavior for a laptop. If you're deploying wireless applications, I recommend familiarizing yourself with the various optimization approaches so you can evaluate applications accordingly.

Using Your Cell Phone as a Modem

Network Computing Mobile Observer, Jun 14, 2006

I have been using cell phones as modems for laptops all the way back to analog cellular over a decade ago. With each phone and with each wireless technology, there has been a learning curve, sometimes a precipitously steep one. But in all cases, I have been happy with the end results. First, let's look at why you might want to do this, and then let's look at the issues with doing so given today's technologies and service plans.

Tethering options nowadays are either via Bluetooth, IR or special-purpose USB cable that connects from your laptop to your phone. In the past, cables also connected to serial ports, but USB is most common nowadays. Tethering is most useful for people who only occasionally use wide-area wireless data and who also do some amount of data on their phones. By tethering, you can use the same data plan for both the phone and the laptop. I find tethering particularly useful in a smartphone/laptop combo when traveling. I use the smartphone for viewing most of my e-mail, responding to a few urgent ones, deleting others, then leaving the rest for my tethered connection, which I may engage in at the end of the day. I'm convinced there are a lot of people running around with laptops and phones who would love to do this, and already have all the right gear, but are either not aware this is an option or, if they do, don't know where to start.

Three developments in recent years have made tethering more practical. One is that an increasing number of phones have Bluetooth capability, which I think is the superior choice for tethering because it's fast enough (though just barely) for today's 3G networks and doesn't require you to carry a cable. The second development is Bluetooth being built into many laptops. The third is a good selection of connection management applications that walk you through the initial configuration and then provide a user interface that shows the connection status and provides buttons for connecting or disconnecting from the network. Some connection managers, which are provided by phone vendors as well as cellular operators, can also manage your Wi-Fi connection. Despite some of the attractions and ready feasibility of tethering, there are a number of complications and issues that have hampered adoption.

One complication with Bluetooth is that not every phone with Bluetooth supports the necessary Bluetooth DUN (dial-up networking) or SPP (serial profile) profile. Even if the device supports Bluetooth DUN, you'll need to make sure the connection manager supports that device. The connection manager is not essential, as you can manually configure a Windows dial-up networking connection specifying a Bluetooth modem, but this borders on rocket science. If you decide to take this route, there are application notes floating around the Internet that provide step-by-step directions for different devices, so do the search rather than figuring it out yourself.

Another complication is that you generally can't use your phone if you're on a data connection. However, your phone will still ring if somebody calls, and if you take the call, your data session will suspend and then resume once your call is over. With new 3G technologies such as Universal Mobile Telecommunications System (UMTS), which Cingular is now deploying, the situation improves because UMTS allows simultaneous voice and data. This means you can be using your cell phone for a voice call even as you are using your phone as a modem for a laptop. In fact, Cingular indicated in a recent interview that the company views its 3G network as particularly well suited for tethering, though the operator also supports PC Card modems and laptops with embedded modems, depending on user preferences.

The final complication is one of pricing plans, where, until recently, operators did not have clear policies. If you're subscribing to a usage-based plan, whether for phones or laptops, things are fairly clear cut. But

where things can get murky is with unlimited plans, which typically run $60 for laptops but anywhere from $10 to $40 for smartphones. Clearly, operators don't want subscribers to sign up for an unlimited smartphone plan and then consume large amounts of data with a laptop. However, operators can monitor how their data accounts are used and presumably clamp down on users abusing their accounts. And if you're already paying $30 to $40 for a wireless e-mail account with an operator, how much extra will you be willing to pay for a tethering plan? Nevertheless, the tide seems to be turning because, recently, all the major operators instituted tethering pricing plans for new devices such as the Blackberry 8700, Palm Treo 700p and Motorola Q.

Add up all the complications and, clearly, tethering is not for everybody. In a recent E-poll conducted by Network Computing to research mobile wireless broadband adoption for a forthcoming feature, only 21 percent of 593 respondents indicated tethering as their preferred approach. Nonetheless, this is still a sizeable segment of the market, and one I expect to grow.

VOIP Over Wide-Area Wireless: A Tricky Proposition

Unstrung, May 25, 2006

For those of you thinking of running VOIP over wide-area broadband wireless services (e.g., EV-DO, HSDPA), here are some things to think about. In the first place, if you read the fine print of the service contract, you likely are not allowed to do so. But secondly, you may not be happy with the results. If you've tried it and it worked, then you were lucky, and were probably running on a relatively unloaded network. I got quite a bit of insight into this topic last week listening to an engineering presentation from Qualcomm Inc. on how VOIP will be implemented for Revision A of EVDO. The presentation was part of a PCCA meeting I chaired last week on the topic of Internet Multimedia Subsystem (IMS).

Sprint Nextel Corp. and Verizon Wireless will be deploying EVDO Rev A in the 2007 time frame. Rev A primarily features a faster uplink, but it also includes provisions such as QOS that will enable VOIP. However, that doesn't mean that the operators will roll out VOIP right away. That will take another year or two as there are a lot of additional items required to make VOIP work.

What Qualcomm aptly showed was that to deliver high-quality, low-latency voice in IP at the equivalent (or slightly higher) spectral efficiency of current circuit-switched approaches is going to require extremely sophisticated communications protocols and voice processing. First there is packet header compression, as otherwise you're sending 20 bytes of IP header information for every 22 bytes of VOIP payload. Robust Header compression knocks this down to 4 bytes. Then there is the elimination of PPP framing overhead, QOS implementation, de-jitter mechanisms, and items called smart blanking and time warping to recover from low-level bit loss.

Qualcomm and other vendors are actively working on these areas. Until all these capabilities are baked into the broadband wireless technology itself, any VOIP usage will be of significantly lower quality, and will consume far more bandwidth than existing voice services. Similar efforts will be required to make

VOIP a reality for technologies such as UMTS and WiMax. It's going to be the end of the decade before you see widespread VOIP over wide-area wireless networks.

Convergence Update: IMS and UMA

Network Computing Mobile Observer, May 24, 2006

There is currently a lot of discussion in the industry about convergence, which is somewhat amusing because I thought we were done with convergence. The first time this buzz word came in vogue was about 10 years ago, with the rise of the Internet, especially as we were starting to inflate the big telecom and Internet bubble that drove the NASDAQ index to stratospheric heights. Well, convergence of all kinds of services, such as voice and multimedia over IP networks, never quite happened back then, despite the efforts of large companies and valiant startups alike. Now convergence is back, and this time it may just stick. I obtained considerable insights into this area from chairing a Portable Computer and Communications Association (PCCA) symposium last week on the IP Multimedia Subsystem (IMS), a convergence platform if there ever was one. The meeting was hosted by Tatara Systems and had in-depth technical content presented by Airvana, Intel, MCCI, Nokia, NMS Communications, Qualcomm and Tatara.

First, we need to clear up the term "convergence," because it actually refers to multiple things. To begin with, there is industry convergence, as companies both consolidate (e.g., wireless operators) and cross into each other's turf (e.g., Google doing metro Wi-Fi). Then thereis device convergence, with devices supporting multiple functions such as voice, video, cameras, e-mail, browsers, gaming and media players. There is also network convergence, which is the notion of people and devices using a combination of wireless and wireless networks. Wireless network options include cellular, Wi-Fi, Bluetooth, UltraWideband, digital mobile TV broadcast, near-field communications and RFID. Finally, there is application and service convergence, where operators can centrally host services such as voice (circuit-switched or VoIP), push-to-talk, messaging, gaming, entertainment and video sharing as well as provide access to these common services regardless of network, be it cellular, Wi-Fi, DSL, coax cable or fiber to the home. IMS, in fact, is a platform for service/application convergence, which is why it is being evaluated and trialed by many large operators. However, no IMS services are deployed quite yet.

We currently have a divergence of convergence approaches. IMS is the favored long-term approach. But there is also Unlicensed Mobile Access (UMA), which is a way for GSM/UMTS voice and data services to be made available over Wi-Fi networks. In effect, UMA handles a small subset of what IMS can deliver. Consequently, it is much simpler. This is a good thing, because if you look at the IMS network reference diagram you'll be impressed by the complexity of the system, with all of its interfaces, gateways and protocols. UMA devices are already available from vendors like Nokia, and they consist of mobile phones with Wi-Fi, cellular (GSM or UMTS) and UMA client software. Although no UMA services are available quite yet, sufficient vendor commitment and operator trials mean that UMA-based services are likely to become available this year.

Why this is important is that UMA (and eventually IMS) will enable a more seamless integration of Wi-Fi and cellular. This means that you'll be able to receive cellular calls with a device operating on your Wi-Fi network. Seamless handover between the two networks will also be possible. Sure, cellular voice minutes are inexpensive in the United States, but the bigger problems this network convergence supports are coverage and, to a lesser extent, capacity. Many buildings now have good Wi-Fi coverage but poor cellular coverage, especially deep inside the building. As an IT manager looking at how to extend voice throughout your enterprise, it's not too early to include UMA on your road map of technologies to monitor. Meanwhile, the converged approach can also add to capacity, which can help in places like airports and shopping malls.

Beyond Wi-Fi/cellular and fixed/mobile convergence, many of the long-term benefits of IMS will be for operators. But the application that we're likely to see in the short term, and that should be of interest to consumers and businesses alike, is video sharing. This allows users with phones that support the feature to be on a call and, midstream, to enable video sharing. The voice call remains circuit-switched (for now), while the video travels on a simultaneous parallel data channel, probably with a slight time lag. Not only might this be useful for showing your grandma your son's antics ("isn't he cute"), it could have uses in tech support ("that yellow cable plugs into the third socket from the right"), real estate ("you'll love this house"), emergency services ("that's the license plate"), auctions ("here's the engine") and so forth. I don't have any hard dates on when video sharing will become available. However, enough of the pieces are in place that it could happen this year, and definitely by next year. These are all baby steps for convergence, but
they are real -- and they will have an impact.

Palm OS Still Alive

Unstrung, May 19, 2006

After Palm Inc. ’s release of the Treo 700w using Microsoft Corp. (Nasdaq: MSFT - message board)’s Window Mobile operating system, many people were wondering what would happen to the Palm OS operating system. The fate of Palm OS was also clouded by PalmSource’s decision last year to discontinue work on Cobalt -- the high-end, multi-threaded OS that was to take Palm into the future. PalmSource indicated at that time its intention to use Linux as its high-end OS platform.

The Treo smartphone line has sold extremely well, with some 3 million Treo 600s and 650s in the market. And it’s no wonder. The Treo 650 does a lot of things well, especially from a useability point of view. So it was no surprise when Palm announced the 700p, an enhanced Palm OS Treo that will support CDMA2000 EV-DO networks. The gadget will address one of the 650’s biggest weaknesses -- namely, its limited memory. The 700p will have 128 Mbytes of memory. I’m sure Palm will sell plenty of these. A UMTS/HSDPA version seems like a logical development as well.

However, while Palm OS has its dedicated following, and is well suited for light-weight applications such as organizer functions and email, it simply does not have the horsepower of competing systems like Windows Mobile 5, which is far more sophisticated in its multi-tasking and networking capabilities. Mitigating this limitation is that most users do not do much beyond organizer tasks, email, and some browsing functions on their smartphones. So it makes complete sense for Palm OS to keep being used for the moment. However, as enterprises slowly figure out how to deploy more demanding applications for smartphones, its likely that many of these will be targeted for alternative smartphone platforms.

Hotspot Redux

Unstrung, May 16, 2006

Lost in all the discussion and arguments about metro WiFi has been the deployment and usage of wireless hotspots. Remember those? Five years ago, working with Datacomm Research Co. , I published the first comprehensive market report on wireless hotspots. While we were enthusiastic about the potential of the technology, particularly in its ability to deliver broadband using relatively low-cost infrastructure, we were concerned about the ability to make hotspots a viable business unless combined with other services such as cellular.

Since then, although wireless hotspots have not taken over the mobile broadband space, the number of hotspots around the world has grown at a strong pace. According to JiWire Inc. statistics for May, there are now 114,150 hotspots in 126 countries. The number is more than double that of a year ago, which makes for impressive growth. The U.S., at 38,588, has the largest number, more than twice the next country, which is the U.K. at 16,898. Pubs, cafes, and restaurants have the lion's share of deployment with 40 percent among them, followed by hotels at 26 percent of locations, then stores at 13 percent. And despite all predictions of free service dominating this segment, according to the stats, less than 10 percent of hotspots today are free.

Even though the market is highly fragmented, and there are no official standards for roaming, the key players in this space have done a great job of providing single-account access to multiple locations. For example, iPass Inc. (Nasdaq: IPAS - message board) boasts access to 50,000 hotspots, Boingo Wireless Inc. to 30,000, and T-Mobile USA has 6,914 of its own hotspots, with access to a total of 22,799 hotspots through roaming arrangements.

We're still a long way from having public WiFi everywhere. It would take about a million hotspots to cover every logical location in the U.S. if you count all hotels, shopping areas, airports, restaurants, and so forth. Nevertheless, in the last five years we've come a long way.

Mobile Broadband - Trying to Catch a Fast-Moving Target

Network Computing Mobile Observer, May 3, 2006

This April and May I've been teaching a series of one-day seminars on the evolution of Wi-Fi, 3G and WiMAX. Along with forcing me to update much of my slideware, these classes have made me assess the state of the wireless-data industry. Although wide-area wireless data revenues are increasing, they still only represent about 10 percent of operator revenues.

What's the holdup? Is what the market really wants technologically achievable? My view is that wireless networks are successful to the extent that they can act as enhancements and extensions to our existing infrastructure networks. However, one fundamental challenge is that the performance of wireline infrastructure increases very quickly. Even though wireless network performance is increasing as well, wide-area wireless networks are playing catch-up against a fast-moving target. The implications are significant, especially for the evolution of the mobile broadband market. Let's look at some numbers.

Close to 1990, when a fast dial-up modem ran at a throughput rate of 9600 bps, I examined the specifications for a forthcoming data service for digital cellular, also at about 9600 bps. My naive view back then was that if a wireless network could offer a modem service as fast as wireline, we'd all go straight to mobile heaven. Alas, by the time the wireless service became available several years later, modem throughput rates had gone up to 28 kbps. Wireless connectivity was sluggish in comparison because of both the lower throughput rates and the higher network latencies. This trend has continued.

There was considerable excitement in the wireless industry as 1xRTT and EDGE networks were being primed for deployment. Finally, wireless would catch up! Yet by the time these technologies became available, with around 100 kbps throughput performance, most serious network users had moved to 500 kbps or higher broadband networks. Dial-up use, meanwhile, was declining as travelers started using services such as hotel-provided broadband.

Now, operators are deploying 3G networks with rates in the 500 kbps to 1 Mbps range. Have we finally caught up? It all depends on your point of comparison. Today's 3G data services are blazingly fast compared to prior wireless services. Yet compared to, say, a DSL connection, many applications operate noticeably slower.

Enhanced versions of 3G will be able to boost throughput rates by an order of magnitude too, approaching 10 Mbps of real throughput, as will mobile WiMAX. But new cable modem and DSL services are already approaching the 5 Mbps range, and with fiber to the business or the home, Internet speeds of 100 Mbps-plus become feasible. So the catch-up game will continue.

Relative to wide-area networks, Wi-Fi has done better on the speed front because it can harness more radio bandwidth per user; it also uses much simpler protocols. In fact, early IEEE 802.11n-like products are beating out 100 Mbps Ethernet in throughput tests. But Wi-Fi is a local-area technology, and despite all the buzz about metro Wi-Fi, it is ill-equipped to become the personal broadband network available anywhere.

The other wireless limitation is capacity. There simply is not enough spectrum available, even with the most spectrally efficient technologies available, to deliver to large numbers of users the gigabytes per month that many Internet users are starting to consume with applications like movie downloads. Thankfully, business users are not generally bandwidth hogs, and 3G networks should be able to satisfy their needs for quite some time.

At the end of the day, wire and fiber will win the speed and capacity battles in most populated areas. Despite this, we must realize that today's wireless networks are still immensely useful and that they will be ever more useful. The RIM BlackBerry has proved that you don't need much bandwidth to provide an extremely useful service. But this is a highly optimized solution involving gateways and a network operations center. What does it take to get beyond e-mail on a widespread basis? Do we have to do it one application at a time, as with e-mail, or are we reaching the point where the networks are fast enough and affordable enough that a whole wealth of new applications will become practical? I'm hoping for the latter.

Google Eases the Third Screen Blues

Unstrung, Apr 26, 2006

Recently Google changed its search page for mobile devices so that if you click on one of the search results, Google automatically formats the subsequent page for better viewing and faster loading on a mobile screen.

This addresses the problem of large and complex Web pages that take a long time to render on a phone or PDA, and that are difficult to read despite the valiant efforts of today's mobile browsers. I've used this Google capability for a while and have found the results to be quite helpful. Unfortunately, there is no way to make this rendering process perfect, because there is no algorithmic process that can extract the essential elements of a complicated Web page for a small screen.

However, because 99 percent or more of Web content was designed for large screens, the fact that Google is trying to make more of the Web accessible to mobile devices is a good thing, in my opinion. This is especially the case as network speeds are improving, phones are becoming more capable, and we're all trying to figure out how to do more with our smartphones than just email and organizer functions.

There are some concerns about what Google is doing, such as the issue of copyright. Does Google have the right to alter the content? It's fine with me, but then I'm not an attorney.

Another concern is what if the content has already been optimized for a small screen? (Unlikely by and large.) Actually, Google provides another search option, which is to select Mobile Web, in which case it will only report hits for pages that are formatted for mobile devices. This restricts search results to a much smaller universe, but the resulting pages should be more readable.

Overall, I think Google is onto something, and many users out there with browser-equipped phones -- especially larger ones such as BlackBerries, Treos, and Windows Mobile devices -- are likely to benefit from Google's approach.

Do-It-Yourself Wireless Email

Unstrung, Apr 18, 2006

Although there are a number of good wireless email solutions available from the likes of Research In Motion Ltd. (RIM) (Nasdaq: RIMM - message board; Toronto: RIM), Good Technology Inc. , Intellisync Corp. (Nasdaq: SYNC - message board), Seven Networks Inc. , Sybase Inc. , and now Microsoft Corp. (Nasdaq: MSFT - message board), there are also increasing options using Internet-based services. As an exercise, I decided to see if I could build my own wireless push email solution for my Palm Treo 650.

Why "push"? Couldn't I just let the Treo poll my POP3 Internet-hosted account every 15 minutes? Well, polling does not quite provide the real-time email experience that people sometimes desire.

Here is how I did it:

First, I created a free Yahoo account. Then I arranged, via a Web page configuration option provided by my ISP for a copy of my POP3 Internet email to be forwarded to my Yahoo email account. On my Yahoo account I configured SMS notification to be sent to my Treo 650 whenever I received an email. Then on my Treo 650 I installed a free utility called "TreoHelper" in which I enabled an SMS trigger that searches for my email address in every SMS message and if it finds it, the trigger launches the Treo VersaMail email client, which then automatically downloads my email from my POP3 email account.

So that I don't have to actually view each SMS message, I enabled an option in TreoHelper to delete the SMS after its VersaMail launch operation. The only catch now is the volume of SMS messages. At approximately 1,000 messages per month, I would need an SMS bucket plan to bring down the SMS cost to $10/month -- otherwise I'd see a $100 SMS charge.

The solution works, though I'm still waffling on using this versus the 15 minute polling approach, as 15 minutes is almost close enough to "real time" for me. Another consideration is that polling uses a little more data and battery. Regardless, the moral of the story is that there are increasing mobile options available from Internet solutions providers to accomplish what used to require dedicated wireless infrastructure.

Integrated Voice/Data in 3G

Network Computing Mobile Observer, Apr 12, 2006

One of the promises of 3G cellular has been integrated voice and data. In this column I examine what this means, why and where it's important, to what extent the different cellular technologies support it, and how this capability will evolve over time. I'm bringing it up now because Cingular is currently rolling out an enhanced version of Universal Mobile Telecommunications System (UMTS) that includes High Speed Downlink Packet Access (HSDPA), which supports integrated voice and data. Meanwhile, Sprint and Verizon have recently made announcements about an upgrade to CDMA 2000 Evolved Data Optimized (EV-DO), called Revision A, that provides a key building block for integrated voice and data, a capability currently not available for EV-DO.

Integrated voice and data means that while you're on a voice call you can simultaneously engage in a data session using the same device. This allows you to do things like send pictures while on a call or do simultaneous voice and video. Not only are there personal scenarios where this might be entertaining, but this ability should prove invaluable for many job situations. Examples include insurance adjusters at the scene of damaged property communicating with headquarters, real-estate agents giving real-time interactive video tours of houses to their clients and medical technicians showing pictures/videos of people they are treating while communicating with doctors. These examples assume a handset device. In a laptop scenario, it depends on your connection. If you're using a PC Card modem or a laptop with integrated 3G capability, then you're likely in a data-only mode, where the integrated voice/data capability is not relevant. But if you're using a phone as a modem, perhaps with a Bluetooth connection, then you can make a voice call in the midst of a data session. And for those of you wanting to use a VoIP system like Skype over 3G, today's networks are simply not up to the task yet.

So which technology can do what? The most capable 3G technology in this area is UMTS/HSDPA, where handset devices support simultaneous voice and data. This is possible because the UMTS voice channels and HSDPA data channels co-exist in the same radio channel. However, in the case of EV-DO, the radio channel is dedicated to high-speed data functions while voice travels over 1XRTT (one carrier radio transmission technology) channels. The 1XRTT channels also support lower speed data services but require devices to be in either a data mode or voice mode.

The lack of integrated voice/data puts EV-DO at a competitive disadvantage relative to UMTS/HSDPA. EV-DO operators have another problem, too. With UMTS/HSDPA, the entire radio channel is available for any combination of voice/data loading. With EV-DO, operators have much less flexibility to dynamically allocate radio resources between voice vs. high-speed data. So far this has not been an issue, as EV-DO networks have relatively few users, but as data services become more popular, operators may have to add EV-DO channels; this assumes, however, available spectrum, and it also takes away from voice capacity. The answer to this dilemma is EV-DO Rev A, which operators will begin to deploy next year. EV-DO incorporates improvements that will enable VoIP functionality, including QoS support, reduced latency and improved uplink capacity. VoIP will enable integrated voice and data for EV-DO, and it will allow more flexible radio resource management, because EV-DO channels now will be able to support both voice and data functions.

However, implementing VoIP over 3G will still be a hugely complicated task. Current approaches of using dedicated CDMA channels for voice are extremely mature. Replicating the same level of voice quality and consistency in the packet domain will involve a lot of radio engineering as well as additional network enhancements such as sophisticated IP packet header compression and an IP core for voice. An important potential element of this core will be the IP Multimedia Subsystem (IMS), a set of infrastructure elements that integrates voice, data and multimedia. IMS is based on the Session Initiation Protocol (SIP) and allows for common core services (voice, voicemail, video, presence, push-to-talk and so on) to be implemented in applications that operate independently of the access network. IMS is under consideration for GSM/UMTS, CDMA2000 and WiMAX networks as well as for wireline networks where it can facilitate wireless/fixed convergence.

I expect the path to VoIP to be a two-stage process for EV-DO carriers. The first step will be deploying the EV-DO Rev A upgrade in 2007, initially as a data-only enhancement, mostly for improved uplink speed and reduced latency. The second part will be installing the ancillary VoIP functions, which probably will not be ready for primetime until 2009. The end game is attractive, especially as it may result in additional voice capacity. VoIP is on the road map for UMTS/HSDPA as well, but the feeling of urgency in this community is much lower because it already has the benefit of a combined voice and data offering.

Pay-As-You-Go Mobile Broadband

Unstrung, Apr 8, 2006

Verizon Wireless announced on March 29 a pay-as-you-go option for people buying notebooks with integrated EV-DO capability. The cost is $15 for a 24-hour period. This allows the operator to compete on a pricing basis with wireless hotspots and other temporary broadband options such as at hotels.

I applaud this pricing option and believe it will be quite popular. There are many professionals, myself included, who don't travel enough to justify the $60 monthly unlimited plans that are currently the norm for 3G wireless-data services. Even though the new Verizon plan only takes four days to reach $60, and realizing that $60/month is only $2/day, not having the recurring monthly expense, and having the freedom to pay for the service as you need it is attractive.

With some hotels charging about $10/day for their broadband services, and others offering it for free, the service only makes sense for people wanting connectivity from a wide variety of locations. And the emphasis needs to be on "wide," because WiFi hotspot service from companies like T-Mobile USA is only $9.99 per day, and, all other things being equal, I prefer a WiFi connection over a wide-area connection from a performance standpoint.

I am curious, however, why the service is only available for notebook computers with integrated EV-DO and not for users with EV-DO PC card modems. Perhaps it's an interim approach for logistical reasons, or perhaps it's to encourage sales of the laptops with integrated EV-DO capability.

Incidentally, I'm a big fan of integrated 3G because of generally better radio performance, ease of setup, and ease of use. Over time, because the pay-as-you-go type of pricing plan simply makes good sense, I'm relatively confident that it will broaden to other types of devices and that it will be adopted by other operators.

Mobile Middleware in the Broadband Era

Network Computing Mobile Observer, Mar 22, 2006

With mobile broadband networks like EV-DO and HSDPA becoming more widely available, many IT managers are asking about the role of mobile (or wireless) middleware. Is it still needed with the higher speeds? The answer is yes, but the reasons have changed.

For background, mobile middleware is the system that operates between a mobile device and a server. Usually it consists of software on a mobile device and a mobile server that acts as a proxy for the mobile node so it can communication with an end server such as a mail server, Web server or database server. The objective of this architecture is to optimize the communication between the mobile node and the mobile server for wireless connections, addressing such unique wireless communications challenges as slower speeds, higher and variable network latency, intermittent connections, changes in IP address because of connections being reestablished and so forth.

In its earliest forms, mobile middleware made it feasible to use really slow networks such as Mobitex and CDPD with applications that would have performed much too slowly in their native form. Today, with typical speeds of 3G networks approaching a megabit per second, speed is less of an issue. In fact, many networking applications will function quite well with no middleware assistance, especially with light use. However, the mobile environment is still fundamentally different from the high-speed Ethernet world over which many networking applications function.

I can see this with my use of a Treo 650 and POP e-mail, where about 20 percent of e-mail downloads fail because of variable delays and occasional slow throughput that confuse either the TCP or the POP3 protocols and result in timeouts. Compare this with a RIM BlackBerry architecture that is rock solid in its mail delivery, regardless of wireless network type. Why? Because it employs a mobile middleware architecture, including the BlackBerry Enterprise Server and RIM's NOC.

The same goes for the other wireless messaging gateways. With sufficient mobile design emphasis, there's no reason that direct client/server interaction cannot be improved for the mobile environment -- as Microsoft has done with Windows Mobile 5 and SP2 for Microsoft Exchange. However, this is not the case for most applications. So we are left with a situation where even Web browsing over EV-DO and HSDPA can be enhanced by a factor of two to four using Web optimization (another form of mobile middleware). With today's busy Web pages, downloading a Web page in 3 seconds instead of 10 over a 3G connection is a noticeable gain.

However, speed is just the tip of the iceberg. Another crucial element is security. Most enterprises employ VPNs for remote access, including 3G wireless use. One approach is to use a standard VPN from the likes of Checkpoint, Nortel or Cisco, which is certainly feasible. But an increasingly popular alternative is to use mobile middleware with VPN functionality (referred to by some as mobile or wireless VPNs) from companies such as Ecutel, IBM, ipUnplugged, NetMotion Wireless and Padcom. Not only do you get the secure tunnels, but you also get transport optimization through compression. In addition, depending on the vendor, you may get important features such as session maintenance, so if you lose connections -- dropping from coverage, for example -- the mobile middleware maintains your application sessions, allowing them to quickly resume once you are back in coverage. Some mobile middleware solutions also allow users to maintain secure applications sessions even as their network changes, say from 3G to Wi-Fi.

Another increasingly important and relatively new function that mobile middleware can potentially support is policy enforcement. This is where the mobile middleware enforces what applications can access the network with what network connection type. If you only have limited bandwidth because you've roamed into a 2.5 G coverage area, it might be nice if bandwidth-hogging applications such as Windows auto-update were put on hold, letting you download your e-mail efficiently. With so much automatic networking activity occurring on today's laptops, this is a feature I can truly appreciate -- currently, I manually disable virus updates, Windows updates, spyware updates and so forth before using my wireless connection.

And what if your enterprise policy is to always use a VPN for networking, but you find yourself at a wireless hotspot where you need to use a browser outside of the tunnel to enter your hotspot login credentials? Some mobile VPNs can let you do this by temporarily allowing a browser session outside the tunnel.

The benefits of mobile middleware are quite clear, but which specific solution to use can be much less obvious. Feature sets and architectural approaches vary considerably amongst the vendors. The challenge is sorting through all the available approaches; not only are there mobile middleware vendors dedicated to this space, but the mainstream application vendors and VPN vendors are also adding mobility functions to their products.

Bottom line: If you or your users are making light use of wide-area wireless networks, you can probably get by without using mobile middleware. But if you're going to do any heavy lifting, where applications will be doing large numbers of networking transactions every day, where you may have a large number of users in the field and user productivity can only be achieved through applications that work extremely efficiently and reliably, then you seriously need to consider mobile middleware as part of your overall architecture.

Wireless vs. Wireline – It's No Contest

Mobility Loop, Mar 22, 2006

With so much emphasis on wireless networking today, you'd think that wireless was about to displace all wire. That simply won't be the case, not because of laws of economics but due to laws of physics. When you look at speeds and capacity, you have to consider the capacity of fiber versus the capacity of radio. Fiber has a theoretical capacity in the range of 10 to 100 Tbps. That "Tbps" is terabits per second, or 1,000 Gbps. Even if you had the entire lower 10 GHz of radio spectrum available to you, and assumed a whopping 10 bps/Hz through the most advanced radio techniques available (likely breaking Shannon's law in any real world deployment with interference), you'd still only end up with 100 Gbps. So, what we have is the entire useful radio spectrum carrying one percent of the theoretical data capacity fiber. Okay, maybe you can only do 10 Gbps over today's fiber system, but remember, that's just one strand. Want more capacity; add more strands. Now, take into account the tiny sliver of spectrum available to any operator, and the ratio of wireless to wireline capacity becomes even smaller.

Thanks to the telecom boom and bust, a lot of fiber was deployed, with only a small fraction of it in use today. My view is that one role of wireless is to complement the fiber in fundamentally two ways: for mobility and for filling in the gaps.

In the case of mobility, we don't want a physical tether as we relish our ability to communicate from anywhere. So this is a perfect role for wireless, keeping in mind that the larger the cell, the more users there may be to share the spectrum. This results in typically lower average throughput rates for wide-area networks compared to local-area networks. In the case of filling in gaps, this is the other good use. In the US, copper is everywhere, but fiber only currently passes a small percentage of buildings and other target locations like broadband wireless cell sites. So while initiatives such as fiber to the curb (FTTC) are exciting, it will literally take decades to hook-up all the fiber. This is where technologies such as mesh Wi-Fi and WiMAX can help fill in the gaps. I love wireless, and see it as a technology that can transform lives and make companies much more productive, but we need to keep its capabilities in perspective.

IEEE 802.11n – Dramatically Faster

Mobility Loop, Mar 17, 2006

The contention over different technical approaches for IEEE 802.11n appear to have been resolved, and the standard is moving towards completion. The IEEE 802.11n working group just gave its tentative approval to a draft 1.0 standard. Next come comments and votes from the members of the working group. If approved by 75% of members, the proposed standard will be formally submitted in July. However, based on 500 members, numerous expected comments, and the likelihood of a second version of the specification before final ratification, it will be well into 2007 before this high-speed standard is complete. And this truly will be high-speed.

As reported by Glenn Fleishman in Wi-Fi Net News on March 15, according to the CTO of Atheros, peak raw rates will be 600 Mbps with 450 Mbps achievable in the fastest configuration, with 150 Mbps realizable in a 200 Mbps base configuration. Assuming about 25 Mbps of achievable throughput for IEEE 802.11a and IEEE 802.11g, IEEE 802.11n will be six to eighteen times faster than today's standardized Wi-Fi networks, which is a huge increase in speed.

Do we need this speed? Absolutely. Do we need it for all applications? Of course not. For Internet access where the limiting factor is the broadband connection at 1 to 5 Mbps, there will be no benefit. But for computer to computer communications locally, it will be invaluable. I regularly have to disable my Wi-Fi connection and plug in Ethernet, to do large directory synchronizations across computers. Consumer multimedia applications, especially combined with 802.11e QoS will be a big application. For instance, moving HDTV requires 20 Mbps of reliable throughput, which is touch and go with current wireless, and how many homes have just one TV screen? Also, the Multiple Input Multiple Output (MIMO) antenna technology will extend range compared to current standards, which in many cases will be as valuable as increased speed.

But be careful before you rush out and buy "n" products. Until the standard is complete and the Wi-Fi Alliance certifies products for interoperability, it could be a risky proposition buying early versions of products based on early versions of the specifications. Yes, you'll likely get the speed, but despite what vendors may claim, there are no guarantees that pre-standard products will be upgradeable to standardized versions later on. This may be fine for smaller deployments, but potentially problematic for larger ones.

TCP/IP and Wireless

Mobility Loop, Mar 7, 2006

Tuesday, March 07, 2006 Written by Peter Rysavy

In early days of wide-area wireless, using TCP/IP was a risky proposition. The first opportunity to do so was in 1994 with Cellular Digital Packet Data (CDPD), which was the first IP-based wide-area wireless network. With CDPD you could run TCP/IP applications, but the biggest problem was how slowly most of them operated. With higher speed networks like EV-DO from Sprint and Verizon, and HSDPA from Cingular, you'd think we'd be in networking heaven, and on a relative basis, we are. Certainly the throughputs are extremely good with typical quoted rates of 400 to 700 kbps. However, as networks have become faster, they still have not completely solved one problem—latency.

Latency is much better than it has ever been before, but it is still in the 200 to 300 msec range, and has high variability. In some detailed testing I did recently with one high-speed cellular network, latency across about two hundred measurements averaged 290 msec as measured by pings against a test server, but the amount of variation in delay had a standard deviation of 285 msec and the highest delay was 3900 msec. These numbers mean that the timers used in TCP's complex algorithms for detecting

network congestion can become confused, the connection can stall and effects can propagate up to the application layer where operations such as a mail download or file transfer can fail.

So yes, TCP/IP works better over wide-area wireless than ever before, but still not perfectly, and not with the same degree of reliability that you can expect with Wi-Fi or Ethernet. The bottom line is that if you are deploying a productivity application that has to be extremely reliable, you may well need to fine tune that application for wireless, or at least test it thoroughly to know that it will respond in different conditions such as weak signals or network congestion.

Figuring Out Metro Wi-Fi

Network Computing Mobile Observer, Mar 1, 2006

Several weeks ago I chaired a meeting on municipal and mesh Wi-Fi developments. Since I found the meeting quite informative, I thought I would report some items I learned. If this topic seems familiar, it's because I wrote a column about it in December 2005, presenting the pros and cons of this technology. The meeting was held on February 2, 2006, in San Francisco by the Portable Computer and Communications Association and hosted by PCCA member company iPass. Leaders in the industry, including BelAir Networks, Cisco Systems, Motorola and Tropos Networks, presented on industry developments, standardization, deployments and opportunities. Combined with in-depth technical discussion, the meeting provided a fairly comprehensive snapshot on the status of the industry. Here are some interesting observations and conclusions.

Some 400 cities in the world are doing some sort of municipal deployment, so clearly this technology has gained some momentum. However, most deployments are quite recent and in relatively early stages. Mesh capability is essential in making metro-area Wi-Fi feasible. By allowing packets to hop from access point to access point, and only having a subset of access points connect back to centralized infrastructure, deployment costs are significantly reduced. The typical number of hops is between two and three. Each hop can add some 30 to 50 msec of latency. Current mesh solutions are all proprietary, but IEEE 802.11s will standardize mesh networking, allowing for interoperability across vendors. However, this standard is not expected to be finalized until 2008.

What I found especially interesting was the analysis of capacity and scalability. Typical numbers of access points range between 20 and 40 per square mile. Resulting aggregate capacity is in the 20 to 100 Mbps range, a noteworthy figure as it is about an order of magnitude higher than current cellular systems, which makes sense as there is about an order of magnitude more access points than base stations per square mile. If you want to service a large number of broadband users in a dense area, clearly Wi-Fi has a leg up on alternative wide-area technologies, assuming you can afford the required backhaul capacity.

I was also impressed by the mobility capabilities of some of these networks. The Motorola solution, designed to support both public-safety and consumer applications, can handle mobile speeds to 100

mph. Those are some really fast handoffs, and my only question is how quickly you would drive out of the coverage area at that speed. If proponents have their way, however, entire cities could be covered, and that won't be an issue. But that won't necessarily be easy. Though it may sound like a simple matter to install access points on light poles, the logistics are quite complicated: Every township has separate control over these poles, and the moment you're around freeways, the federal government has control of the lights. Network operators must involve a lot of entities to achieve widespread coverage. In addition, some poles are powered at 110 volts whereas others are powered at 440 volts. These challenges are reminiscent of the problems Metricom had rolling out its networks.

The one question that nobody was able to clearly answer at the meeting was how people would manage interference between the mesh and other networks. This remains a complicated problem because of the huge variety of radio circumstances that will arise. To minimize deployment cost, the mesh access points are located in the loosest configuration possible. If you have 25 access points per square mile, that translates to 1,000 feet between access points and relatively weak signals if you're at the maximum distance of 500 feet from the access point. This is why many of these networks call for repeaters on the sides of houses to provide indoor coverage. If you're at the weak part of the signal, your private Wi-Fi network will simply win out over the mesh signal. But if you're at the weak point of your own coverage, such as in a distant bedroom, and there is a mesh node just outside your house, that signal may prevail over yours. You may be able to switch channels if your neighbor is not already using that channel. Or the IEEE 802.11 carrier sense protocols may be sufficient for co-existence if all the appropriate nodes can hear each other. Or maybe mesh networks will stay at 2.4 GHz, where the frequency is needed for propagation, and shorter range private networks will migrate to 5 GHz. How this all shakes out remains to be seen. For public safety, government users have the option of using the licensed 4.9 GHz band, which at least isolates them from unlicensed networks.

Bottom line, I still see significant challenges for wide-area mesh networks, but their deployments have sufficient momentum that it is inevitable that there will be some level of deployment. Current lack of standardization definitely puts municipalities at risk in the event that their vendors go out of business, as has already happened to some cities that were using products from Vivato. Interoperability aside, today's vendors do have highly capable products that can support relatively broad coverage areas. However, while some people might be satisfied with the resulting coverage for fixed or nomadic applications, mobile applications are likely to need greater coverage than offered by the mesh, and so the need for much wider 3G services will not go away. As in my own use of wireless technologies for productivity purposes, I see an increasing number of applications best served by a combination of wireless technologies.

3GPP Long-Term Evolution

Mobility Loop, Feb 21, 2006

Get ready for another wireless acronym – LTE. I suppose it's in the nature of wireless technology that as soon as you start widely deploying one version of a technology, you start discussing the next version.

This serves to generate market excitement and keeps thousands of engineers employed working on the next big thing. Unfortunately, it can make current offerings seem inferior compared to the glowing capabilities of what is coming down the road.

The most current version of deployable 3GPP technology is Universal Mobile Telecommunications System (UMTS) High Speed Downlink Packet Access (HSDPA), as being rolled out by Cingular Wireless and other operators around the world. This is based on Release 5 of Third Generation Partnership Project (3GPP) specifications. Next year we can expect Release 6 features such as High Speed Uplink Packet Access (HSUPA), and then later with Release 7 we'll see items like VoIP. These releases represent the short range roadmap for GSM/UMTS specifications. The longer term roadmap is what LTE is all about. LTE is the next big radio overhaul, with all the best-of-breed techniques available, including OFDM and MIMO. It will offer flexible channelization from 1.25 MHz to 20 MHz, peak theoretical downlink rates to 100 Mbps and latency down to 10 msec in the radio link. If this all sounds very similar to Mobile WiMAX, that's not a coincidence. LTE is intended to match all of WiMAX capabilities, though WiMAX proponents would claim they are further ahead in their efforts.

Why all the attention now? Because key players in the 3G space are positioning themselves to compete in this future wireless space. For example, Nortel announced on Feb 14 at the 3GSM World Congress in Barcelona "As part of the 3rd Generation Partnership Project (3GPP) for the Long Term Evolution (LTE) of wireless network standards, Nortel has achieved approval for the introduction of high speed OFDM and MIMO packet access, in line with its technology HSOPA." HSOPA stands for High Speed OFDM Packet Access. You can expect a lot more LTE positioning like this from other vendors. What does this mean to users in the next couple of years? Not much really, other than an indication of future network capabilities in the 2008 or later time frame. Just keep in mind that there is often a 10 to 1 ratio between the claimed peak rates and typical user rates, and that it always takes longer for new wireless technologies to become available than expected. For all practical planning purposes, stick to what is available now. In the cellular world, that's EDGE/UMTS/HSDPA and 1xRTT/EV-DO.

Sprint Nextel's Jump to the Future

Network Computing Mobile Observer, Feb 8, 2006

Many eyes are fixed on Sprint Nextel as the company contemplates how to take advantage of the large amount of spectrum it owns in the 2.5 GHz band. This is the Broadband Radio Spectrum (BRS), formerly known as Multichannel Multipoint Distribution Service (MMDS). Sprint and Nextel each had sizeable holdings, and the merged company now controls this spectrum in most first-tier markets nationally. We're talking about a sizeable amount of spectrum, close to 100 MHz in many markets, according to one of my associates who tracks these things.

Under FCC rules, the company must use it or lose it, so Sprint Nextel is closely looking at various technologies. Last year, Nextel did trials using Flarion's Flash OFDM technology, which is now owned by Qualcomm. This year, Sprint Nextel will do a WiMAX trial using Samsung's WiBro gear. Developed in

Korea, WiBro is the basis of mobile WiMAX and will become one of the mobile WiMAX profiles. As such, it will be the first "version" of mobile WiMAX available. WiBro service is also being turned on in Korea this year in the 2.3 GHz band. How WiBro performs will be the first good indicator of how well mobile WiMAX will perform and whether it can outperform current 3G technologies, a hotly debated topic.

However, Sprint Nextel is not stopping with WiBro. It is also doing a trial with IP Wireless' Universal Telecommunications System (UMTS) Time Division Duplex (TDD) system. UMTS-TDD is an official 3GPP technology that is essentially High Speed Downlink Packet Access (HSDPA), but with a time division duplex radio. Just last month, Sprint Nextel invested $10 million in IP Wireless, following a $4 million investment last July. This is quite interesting given that there generally is fierce competition between 3GPP (GSM/UMTS family) and 3GPP2 (CDMA2000 family) technologies and that most operators have sided with one technology family or the other.

And to confuse things further, IEEE 802.20, the IEEE standard on mobile broadband wireless access, is now revving up again thanks to Qualcomm's purchase of Flarion and Flarion's Flash OFDM technology. Qualcomm originally saw IEEE 802.20 as a threat to EV-DO, but in my opinion it now sees IEEE 802.20 as a vehicle for competing in next-generation wireless technologies. This way, the company has its bases covered in CDMA, with EV-DO Rev A and Rev B, and in OFDM, with Flash OFDM and IEEE 802.20 standardization. I'm sure Sprint Nextel is monitoring IEEE 802.20 developments too.

Clearly, Sprint Nextel wants to keep all its options open as long as possible so the company can make the best technology decision. Whatever technology it uses, the resulting network, where deployed, will be a high-capacity high-performance network that is likely to outperform Sprint Nextel's current EV-DO network. The company has indicated it will be deciding on the technology this year, making Sprint Nextel the first major wireless operator in the United States to commit to a next-generation wireless technology.

The technology choices are fascinating, but it will be especially interesting to see how Sprint Nextel chooses to market and price the service once it is broadly available in the 2008 time frame. Will the company use it to off-load data from its EV-DO network? You can make a good argument for that, as there is only so much capacity in the Sprint 3G network -- roughly 1.5 Mbps per cell site per EV-DO radio channel, which is not that much. Or will Sprint Nextel use it to compete with DSL and cable? This is a trickier proposition because these users demand far more capacity, anywhere from 300 megabytes per month to 1 gigabyte per month on average based on my calculations, though I will confess that I have found it challenging to get good data. (Contributions on this subject are gladly accepted.) Or will Sprint Nextel emphasize multimedia applications? Opportunity also exists for local bypass of telephone companies, which must be a tempting target, including local phone service and T1 circuits for businesses.

Although I understand the technology options, I'm a lot less clear on what the successful business models are in the wireless broadband area. I do truly believe there is a mass market for mobile broadband, whether delivered to laptops, smartphones or mobile multimedia devices. But when you

factor in the huge expense of deploying a nationwide network, the large amounts of data that multimedia applications can consume, the finite amount of capacity in wireless connections compared to wireline alternatives, and the limited amount of money that users are willing to pay, is there a business model that actually works? A lot of people, including myself, are going to be watching Sprint Nextel closely over the next several years to find out.

Mobile Content – Getting Better Slowly

Mobility Loop, Feb 3, 2006

When I find myself having to wait in places like bank lines, post office lines, or alone at a restaurant, the first thing I do is whip out my Treo smartphone and surf the Web. This usually keeps me occupied for about fifteen minutes, before I've read much of the available content. Then I'm out of luck, because the choices are so limited. Whereas the World Wide Web on a regular computer for all practical purposes contains an infinite amount of content, the amount of entertaining or interesting content formatted for small screens is much more limited.

I usually use Yahoo (http://wap.oa.yahoo.com/) because the site does a fairly good job formatting for small screens, is well organized, and has a reasonable amount of news. In fact, M:Metrics, a company that does mobile-oriented market research, just announced in a press release on January 30 that Yahoo is the leader in this space, with 12.8 million US mobile subscribers accessing the site in December 2005. AOL is second place with about 9 million, MSN and Google following with about 7 million, and ESPN the sports leader with about 5 million.

The operators themselves have portals. I sometimes use Cingular's MEdiaNet site (powered by InfoSpace), but find that it has some rough edges, such as displaying text in only the left half of the screen. Overall, we've come a long way with mobile content, and tens of millions of subscribers is a great start, but in the big picture of things, only ten percent of mobile phone subscribers are accessing mobile content, so we still have a long way to go in both subscribers and in the depth and breadth of the content. Clearly, the level of effort by content providers is commensurate with the amount of access, hence the relatively minor offerings so far. I'm personally looking forward to the day that I can spend a couple of hours reading a whole issue of New York Times. That way I'll be covered no matter how long a wait in line I have.

The Resurrection of IEEE 802.20

Mobility Loop, Jan 25, 2006

Last week, IEEE 802.20 met in Hawaii, with companies making significant technical contributions. All indications are that IEEE 802.20, nicknamed Mobile-Fi, is being brought back from the dead. What's changed? Everything. IEEE 802.20 was initially a vehicle that Flarion Technologies used to try and standardize its Flash OFDM technology. But it never made much standards progress due to lack of

cooperation from other participants who had other agendas, including Qualcomm which at the time viewed Flash OFDM as a competitive threat to CDMA2000 EV-DO.

However, now Qualcomm owns Flarion, and is pursuing opportunities to advance its technology portfolio across all future wireless fronts. This includes EV-DO (which it invented), UMTS/WCDMA (where it is a leading chip supplier), and OFDM-based technologies (where it now has a reasonable intellectual property-rights position with the Flarion acquisition). Qualcomm now is pushing IEEE 802.20. If it actually achieves standardization and gets to market, IEEE 802.20 will compete directly with mobile WiMAX (IEEE 802.16e).

Despite the immense size of the global wireless opportunity, I don't see the market as large enough for three versions of 3G (EV-DO, UMTS and the Chinese TD-SCDMA), the evolved versions of 3G (e.g., 3GPP Long Term Evolution), WiMAX, mesh Wi-Fi and now IEEE 802.20. Something has to give.

From a standards point of view, WiMAX is much further along than IEEE 802.20. However, Flash OFDM is a proven, tested technology, though its not necessarily exactly what is being proposed for IEEE 802.20. (I haven't had a chance to figure out the differences but will try to find out for a future blog.)

The interesting question to me is what effect IEEE 802.20's revival will have on the mobile broadband market. For example, it could prove disruptive to WiMAX if operators decide they need to evaluate one more wireless technology option.

The Smartphone Conundrum

Network Computing Mobile Observer, Jan 18, 2006

Palm's official unveiling of its Windows Mobile Treo smartphone, the Treo 700w, at the beginning of this month made me contemplate the current status of smartphones and their likely evolution. Overall, I'm enthusiastic about this platform, but I have to confess that from a productivity point of view, it's easy to invest just as much time working with these devices as any time you will save from using them. That's the conundrum.

First, let's look at the some of the recent smartphone developments. The first Windows Mobile 5 devices have now appeared. These include the Treo 700w, available from Verizon Wireless and the PPC-6700 from Sprint, both with stylus support and small keyboards. Meanwhile, Cingular started selling its Cingular 2125, which is based on Windows Mobile 5 but more compact -- it uses a regular telephone keypad instead of a stylus and keyboard. There are more Windows Mobile 5 PDAs without phone functionality than there are with phones for the simple reason that adding wireless to these devices is complicated, and wireless operators have rigorous acceptance procedures. Windows Mobile 5 will have a big impact as it addresses some of the serious shortcomings of the prior version, including volatile memory and the need for two-handed operation with a stylus. Anyone who has spent much time with a RIM Blackberry of Palm Treo understands the importance of one-handed operation. The

Treo 700w even shows the battery level, a reading that required multiple steps with the Pocket PC 2003 platform. Meanwhile, connectivity speeds are increasing with RIM Blackberrys, which now support EDGE and EV-DO. RIM also recently released its attractive 8700 line. And PalmSource has decided to use Linux for future platform development. Finally, the Symbian platform gets its best shot for visibility in the United States with the Nokia 9300 being sold through Cingular.

Smartphones are intriguing, as they represent efforts to concentrate the greatest amount of computing and networking power possible in a small form factor that is comfortable to carry. The result is tremendous capability, with a good range of productivity or entertainment applications, such as e-mail, Web browsing, Web applications, file access, portable documents, database access, music, video and hundreds more (just search the online catalogs). For most workers, smartphones can't replace a laptop, but the multitude of features present makes it increasingly possible for many workers to use their smartphone for short trips or other periods away from the office.

However, this level of capability also means complex computer systems and all the associated maintenance and reliability issues. Although I'm reasonably happy using an EDGE-capable Treo 650, I have now updated the system software twice to newer versions as they became available. I was glad each time, because the updates addressed various problems. Although Palm documented the update steps clearly, these were still complex operations that likely defeated many users. I also spent a considerable number of hours getting Bluetooth to work for various usage models, such as using the phone as a modem, for HotSync and with my wireless headset. And still, on a regular basis, the phone will lose awareness of the headset, and I have to pair the two devices again. I've searched various support forums and have found other users with the same problem, so I don't think it's just me. As for stability, it's not bad, but not perfect. My device runs for about a week before I have to reset it. These are just examples of my usage. If you want to see the scope of support issues for these devices in general, just read the thousands of entries in the various support forums for the different platforms.

When I add up my support time, I have to question how much further ahead I am from a productivity point of view. From an IT perspective, these devices must be managed; their software must be kept up-to-date for potentially large numbers of devices, preferably through some centralized system; and remote-access security architectures must be updated to support smartphones, all adding up to considerable time investments. And if you don't make these infrastructure investments, you'll pay even more in compromised security and supporting devices on a case-by-case basis.

Still, I see the ascendancy of the smartphone platform as inevitable. One reason, after functionality, is that prices are dropping, with good smartphones now available in the $200 to $300 range, putting them in reach of a lot of people. Gartner in a July 2005 report predicted global sales of 200 million smartphones by 2008, over 20% of all mobile phones sold. However, unless smartphone vendors want to see an IT backlash against these devices for wasting too much support time, they should make sure their platforms and included applications are as stable as possible. It took quite a few releases of Microsoft Windows before it finally became reasonably stable with Windows 2000 and Windows XP.

After this long painful experience, IT and user tolerance for system instability has gone down, so let's make sure we don't have to repeat this process with smartphones.

Mobile Phone Music

Mobility Loop, Jan 12, 2006

There's no question that people want to carry fewer devices. Witness the struggling PDA market thanks to smartphones that now include PDA functionality. I used to carry a phone and iPAQ Pocket PC, and could even do nifty things like download e-mail to my iPAQ via a Bluetooth connection to my phone that acted as a packet-data modem. But I'm much happier with a single device, currently a Treo 650. Similarly, I don't carry a separate music player but play MP3 songs on my Treo 650, either using headphones or through a Radio Shack tape deck adapter in my car. Though the Treo 650 is less than ideal for music, the concept is powerful, and further evidence that our "phones" are becoming mobile electronic life enhancers.

A lot of companies are now pursuing music on cell phones, though I would say we are still far from musical mobile nirvana. Motorola has their "ROKR" phone, available from Cingular, that is iTunes compatible, but only stores 100 songs. Sprint has a music download service, but songs cost $2.50 each. Most recently, Verizon last week announced at the Consumer Electronics Show their Vcast Music service where users can download songs to a handset at $1.99 each (a little pricey), or 99 cents to a computer. The controversial twist here is that the specific mobile phones supporting this service use Windows Media player. While the phones can play MP3 format songs, once users invoke the service, the MP3 capability is disabled on the phones. However, if a user already has MP3 songs, Windows Media player will transcode the songs to Windows WMA format, so they can be played. We'll see how users respond to this, as even iPods can play MP3 format songs. I believe that any friction like this, similar to cellular's failed attempts in the past with walled-garden content, discourages widespread adoption.

Still, one million available songs and Verizon's marketing push will likely have a big impact. I believe that one way or another, a lot of people are going to sign up for music on their cellphones. Another big enabler that will begin rolling out this year will be Bluetooth stereo headphones that also support phone voice functions.

The Perils of the Bleeding Edge

Mobility Loop, Jan 5, 2006

In wireless networking, there is a fine balance between proprietary systems that provide enhanced capabilities and the safety of standardized solutions supported by multiple vendors. Nowhere is this more evident than in the expanding field of public Wi-Fi systems where vendors are using multiple approaches to extend the reach of Wi-Fi over wide areas, in essence trying to morph a local area technology into a wide area technology. One of the pioneers in this area was Vivato Inc., which used a

beam forming radio technology to dramatically improve range. Unfortunately, at the end of last month Vivato went out of business, leaving a number of municipal Wi-Fi networks built with Vivato's technology asking questions about how to maintain their networks.

The current most popular approach for municipal networks is not Vivato's approach, but to use mesh technology, where access points don't all have to be connected to a backbone. Access points forward packets from other access points towards nodes on the backbone, thus significantly reducing overall deployment costs. Though leading vendors such as Cisco have gotten into the game, helping legitimize the mesh concept, all current solutions are vendor specific, putting municipalities that use these products at risk should their vendors go out of business. Fortunately, mesh Wi-Fi standardization is coming, through IEEE 802.11s. However, it is too early to tell what impact this standard will have, and to what extent risks of vendors going out of business will be mitigated. It is one thing to have a Wi-Fi network in a building become obsolete. It is another matter entirely when the Wi-Fi network covers an entire city.

For those of you wanting to learn more about municipal and mesh Wi-Fi technical developments, I recommend a meeting I'm chairing for the Portable Computer and Communications Association on Feb 2, 2006 in San Francisco that will investigate this topic in detail.

HSDPA Speed: Lab versus Real World

Mobility Loop, Jan 2, 2006

Recently I complained about the inflated speed claims being made about a future version of EV-DO, (EV-DO Revision B). So it's only fair to comment on the misleading speed claims coming out about HSDPA. While I am enthusiastic about the capabilities of both the major 3G technologies, whether EV-DO or UMTS/HSDPA, I think it's important to understand the basis of the claims, and what can actually be expected in the real world. A case in point is Nortel's announcement on December 22 about achieving a 3.6 Mbps HSDPA throughput rate using a card from European-based Option in a test conducted at Nortel's research campus in Chateaufort, France. High Speed Downlink Packet Access (HSDPA) is the version of UMTS that Cingular is currently deploying.

Here's the deal. The 3.6 Mbps number is real, but you as a user will never see it. It is based on using the highest order modulation available (16 Quadrature Amplitude Modulation), and five high-speed codes allocated to the same device. Even then, it is at the link layer, meaning that it does not take into account IP and TCP packet overhead, though this is likely less than 10%. It is effectively the capacity of most of the cell site directed at one user under optimal radio conditions. But even if you were the only user in that cell sector with an excellent signal, while the radio could deliver this throughput, the bottleneck in the real world is the backhaul connection to the base station, which is typically based on T1 circuits that are limited to 1.5 Mbps of throughput. I can guess that Nortel has a fiber connection to their test cell site in Chateaufort, but that's not the real world, though to be fair, some cell sites are on fiber. But most are not.

The reason you'll see actual throughputs in the 400 to 700 kbps, as quoted by Cingular, is two fold. One, as just mentioned, is the real world constraint of backhaul capabilities, which by the way is no different for Wi-Fi hotspots. The other is that multiple users will be active in the cell, sharing the radio capacity. So beware of the inflated speed claims. While they may represent the real capabilities of individual devices, and the theoretical capabilities of the technology, they will not necessarily be what you will experience.

Municipal and Mesh Wi-Fi

Network Computing Mobile Observer, Dec 21, 2005

Despite significant challenges, an increasing number of cities are planning to deploy Wi-Fi networks across large areas. High-profile cities currently include San Francisco, Philadelphia and Houston, but as described in a new report titled "2006 Muni Broadband Outlook" from Light Reading Insider, other cities include Anaheim, Calif.; Atlanta; Austin, Texas; Boston; Buffalo, N.Y.; Chicago; Denver; Kansas City, Mo.; Lenexa, Kan.; Los Angeles; Miami Beach, Fla.; Milwaukee; Minneapolis; Oklahoma City; Phoenix; Pittsburgh; Portland, Ore.; Sacramento, Calif.; St. Paul, Minn.; and Seattle.

There is a considerable amount of debate over whether these deployments are going to become the favored form of broadband wireless access or are doomed to failure. It's a good question, as a case can be made for both sides. That's one reason why the Portable Computer and Communications Association, where I chair the standards and architecture committee, will be holding its next meeting on the topic of municipal and mesh Wi-Fi. We will have technically oriented presentations from the leaders in the field, including BelAir Networks, Cisco, Motorola and Tropos Networks, and good representation from the computer, networking and wireless industries. The meeting is on Feb. 2, 2006, in San Francisco, hosted by iPass, and is open to anybody who wishes to attend, though there is a meeting fee for non-member companies.

Allow me to muse on the cases for and against. The biggest case "for" is one of capacity, as I've already written about in various reports. It's a matter of trivial mathematics that sharing spectrum over a small geographic area versus over a far larger cell site results in a much larger capacity per user. This is because that same spectrum can be re-used over and over. The result is that Wi-Fi can deliver some two orders of magnitude greater bits per second per square meter than cellular data networks. Simply consider a one-kilometer radius cell site versus a 100-meter radius Wi-Fi access point, and that's a hundred-to-one difference in coverage area.

Because of the simpler CSMA protocols and smaller networks, Wi-Fi also has inherently lower latency, less than 10 msec, whereas 3G networks are struggling to get down to 100 msec. As a result, Wi-Fi provides a better broadband experience and scales better for dense user configurations. Another pro is the prevalence of Wi-Fi-equipped computers, whether notebooks or handhelds. The architecture of mesh makes a lot of sense, too, namely that only a subset of access points need a connection to a

wireline infrastructure. With standardization coming through IEEE 802.11s, we'll eventually have interoperable mesh infrastructure. And with VoIP becoming a reality for Wi-Fi, even voice service becomes a possibility.

On the "against" side of the equation, the primary issue is one of spectrum management. For reasonable coverage, mesh deployments favor the 2.4 GHz band using IEEE 802.11b or IEEE 802.11g. However, there are only three non-overlapping channels in this band (1, 6, 11). One common approach to continuous coverage of an area is to use all three channels, effectively using up the whole 2.4-GHz band. Any privately deployed network, or other public network, will have to co-exist in the same radio channels. The CSMA protocols allow for this to some extent, but only if all the nodes involved can hear each other--not always the case for nodes inside a building. How bad is the problem? It's hard to say.

In discussing this with my colleague Dave Molta, his observation is, "If the mesh router is in close proximity (operating at maximum-allowed transmit power), you're likely to experience significant co-channel interference. However, if the signal level has attenuated significantly, as is often the case (in fact, that's one of the big problems with mesh deployments that try to get acceptable in-building penetration from externally mounted routers), then your local signal will overpower the signal from the mesh router. You may still experience some level of co-channel interference, but local performance will probably be adequate." How big of an issue interference is remains to be seen, but we'll find out soon enough.

Then there are the logistics of deploying a huge number of access points over a wide geographic area as well as managing and maintaining them. There will also inevitably be holes in coverage, especially indoors. Whether these systems can be made reliable enough for critical applications such as public safety, other than in select areas, is doubtful. As for VoIP, this typically requires an additional 10 dB SNR compared to data, thereby requiring even more access points.

Bottom line: I see Wi-Fi becoming increasingly prevalent in public areas. Given the unlicensed spectrum and multiple entities using the spectrum as well as the inherent problem of providing continuous coverage with low-powered transmitters, however, I don't see these networks spreading everywhere. Rather, I see them as providing patchworks of coverage in select areas like downtown and in some communities. Nor will they replace 3G data capabilities, which will always be able to provide much better continuous coverage over much larger areas, though at somewhat slower throughput rates. I continue to see public Wi-Fi as highly complementary with 3G. We are quickly moving to a world of multiple wireless access networks. Knowing which to use for what applications will be the tricky part.

Ultra Wideband Standardization

Mobility Loop, Dec 15, 2005

I've always been a big fan of Bluetooth, though I must admit I have wasted more hours getting Bluetooth connections to work between my various peripherals than with any other networking

technology. Currently I use Bluetooth for syncing my Treo 650 to my PC, for using my Treo as a modem for my laptop and for connecting my Bluetooth headset. It's quite evident that personal area wireless makes a great deal of sense, and so I anticipate good things for Ultra Wideband (UWB) over the second half of this decade with its blistering rates up to 1 Gbps. That way I'll be able to sync my PDA in one second. Actually, better uses will be to connect all the high speed, multimedia peripherals working their way into our digital homes; for example no more messy cable going to the high definition TV.

Unfortunately UWB has been deadlocked in IEEE 802.15.3 standardization efforts, with competing Multiband OFDM and direct sequence approaches unable to reconcile with each other. The OFDM approach is championed through the WiMedia Alliance and backed by a lot of companies, including HP, Intel, Kodak, Microsoft, Nokia, Philips, Samsung, Sony and TI among others.

It was notable when the WiMedia Alliance announced on December 8, 2005 that it had been able to get its Physical and Media Access Layers approved as UWB standards through the European based industry association Ecma International. Ecma is a non-profit organization that has issued standards in near-field communications, DVD interchange standards, scripting languages, communications technology and consumer electronics. To download the new standards, go to this site and this one.

Ecma standardization does not carry the "weight" of IEEE standardization, but given that IEEE standardization for UWB appears to be going nowhere, Ecma standardization definitely gives MB-OFDM UWB legitimacy. You can expect products to start rolling out in 2006. For higher-level interfaces, MB-OFDM UWB will use USB, IEEE 1394 and even Bluetooth, facilitating product implementation using existing driver and application architectures. And if some of the Bluetooth type of issues such as managing pairing relationships are properly addressed (a big if), UWB should have a bright future.

Mobile Broadcast Technology Developments

Mobility Loop, Dec 7, 2005

Verizon Wireless last week announced it is going to be the first major carrier to support QUALCOMM’s MediaFLO technology. This announcement is indicative of some important developments in the wide-area wireless space, namely the deployment of mobile broadcast technologies to complement 3G and evolved 3G networks. Though services such as V Cast from Verizon and equivalent services from other operators provide distribution of video and music already, there is an inherent disconnect between the amount of bandwidth that multimedia can consume (particularly with video content) and the limited capacity of cellular-data networks. It is far more efficient to have a single stream of data received by multiple subscribers than separate data streams to each individual user. Hence all the activity. The relative capacities of these networks is an item I studied for a recent report titled “Hard Numbers and Experts’ Insights on Migration to Evolved 3G and 4G Wireless.”

QUALCOMM’s MediaFLO technology is based on OFDM, with FLO standing for Forward Link Only. QUALCOMM is actually building out its own national network that cellular operators can contract with,

using former UHF channel 55 in the 700 MHz band. Devices will need separate radios to access both the 3G and FLO networks. The network is expected to deliver up to 20 live streaming video channels and 10 audio channels, and also some short on-demand programming.

There are a variety of other broadcast and multicast technologies in the works. There is Digital Video Broadcasting Terrestrial (DVB-T), for which a number of Mobile TV networks are being deployed globally, and a mobile-optimized version called Digital Video Broadcasting-Handheld (DVB-H). There is also EV-DO Platinum Multicast which allows multicast slots within EV-DO channels, and Multimedia Broadcast Multicast Service (MBMS), which will allow multicast and broadcast channels within UMTS/HSDPA.

While it's not entirely clear how many subscribers will be willing to pay how much for these emerging multimedia services, it is clear there will be no shortage of wireless technologies to support them.

Linksys WRT54G -- Update

Network Computing Mobile Observer, Dec 7, 2005

Since our story last week about Cisco/Linksys changing the WRT54G Wireless-G Broadband Router from a Linux-based system to a WxWorks-based system, Linksys contacted us to say it will maintain the Linux-based system under a new model number of WRT54GL. The company's statement was, "We are honored by the overwhelming loyalty of Linux users to our products, and this dedicated SKU is our way of saying thank you for the support this community has given to us and our products."

We are pleased that the Linux-based version (current equivalent to WRT54G v4) will remain available for those who desire it. However, we're still not happy about the apparent quality and performance issues of the currently shipping WRT54G. Also, it's a little unusual to create a new SKU for an existing product and to use the old SKU for, effectively, a new version of the product. However, it does mean that Linksys/Cisco can direct its huge volumes of this product's sales, estimated at hundreds of thousands per month, to the version that costs the company less to manufacture. Hopefully, Linksys/Cisco can quickly address the problems with this version so that these hundreds of thousands of users don't have the unfortunate experience I had.

Linksys WRT54G -- Not What It Used to Be

Network Computing Mobile Observer, Nov 30, 2005

The Cisco Linksys WRT54G is one of the most popular SOHO 802.11 b/g wireless routers available. However, because of a recent redesign, this product is nowhere near as reliable as it used to be. Worse yet, it's difficult to tell whether the version you are getting is the old version or the new version.

I stumbled across this a couple of weeks ago when I decided to update my home small-office router, a Linksys BESFR41 v1. Though functional, I had to reset the product once or twice a week, and having had good success with a WRT54G (router, four-port Ethernet switch, 802.11b/g access point) at another location, I decided to get one for my office. As part of my due diligence, I contacted my local ISP, which has well-informed technical support, and asked them about the WRT54G. They assured me it was "rock solid." That was good enough for me.

After obtaining the product, a WRT54G v5, I installed it and instantly found my first problem: extremely slow Web admin screen updates that took five to 10 seconds to update. I then looked at the firmware version: 1.00.0. That was scary. Then I downloaded the latest version of the firmware, release 1.00.2, dated Nov. 1, 2005. This fixed the slow admin page updates, and I completed the installation. It was then that the real problems began. The router ran fine for a while, then I lost wireless connectivity but still had wireline connectivity. I checked all the router wireless settings and found that the MAC address access list I had enabled had disappeared. A one-time glitch, I thought, so I reentered the information. Things continued to work for another several hours. Then I lost Internet connectivity again. I reset the router, which restored operation. However, from that point forward, I had to reset the router at least a couple of times a day. Another problem occurred when I tried to change wireless security settings. I was unable to do so without first restoring the router to its factory settings--not exactly convenient and anything but rock solid.

I then started researching these problems on the Internet and found no shortage of complaints about the product, including stability and throughput issues. The comments were all consistent: people complained that their WRT54G v5 products were unstable, whereas their WRT54G v4 products were highly stable.

Further research into this issue revealed that Linksys changed the operating system in v5 from Linux to VxWorks. Motivations for this change are not entirely clear, though the open-source license requirements for Linux did require that Linksys publish the source for the router. This resulted in people customizing software for the router, thereby allowing the device to perform the functions of repeaters, bridges, VPN servers, VoIP gateways, etc. Linksys does not support any of these extensions, of course, but the platform did create a legion of devoted supporters and developers. It's likely that scrutiny of the source also helped debug the software and may be one reason the earlier versions are as stable as they are. By going to VxWorks, Linksys closes off and retains full control of the platform.

However, the two versions are almost indistinguishable--same model number, same case, same packaging and same screens. The only difference is that the model number on the bottom of the unit indicates the version number; the serial numbers for v4 begin with CDFA and the serial numbers of v5 begin with CDFB. Luckily, these are on the outside of the package. Since my local Wal-Mart carried the WRT54G v4, I eagerly replaced the router, and for the last 10 days it has indeed been reliable, with not a single glitch of any kind.

I personally feel that Linksys has executed a "bait and switch" by not revealing the change, keeping the model number the same and not even putting the version number on the package. I contacted Linksys a week prior to writing this column to offer the company a chance to comment but received no reply. I would like to hope that Linksys will have the commitment to bring the v5 quality level to that of v4, but having no assurance of this, I went back to Wal-Mart and bought a spare v4 router. It was only $50. Maybe that was Linksys' plan all along--to get paranoid owners to double up on their router purchases!

EV-DO Revision B: Speed Hype

Mobility Loop, Nov 23, 2005

The CDMA Development Group (CDG) on November 16 issued a press release about CDMA2000 1x EV-DO Revision B, indicating that the standard would be published in the first quarter of 2006, and stating that “Rev. B introduces a 64-Quadrature Amplitude Modulation scheme, and will deliver peak rates of 73.5 megabits per second in the downlink and 27 Mbps in the uplink through the aggregation of 15 1.25 MHz carriers within 20 MHz of bandwidth.”

Deployment of EV-DO Rev B could happen by the 2008 time frame. It's not a coincidence that Mobile WiMAX deployment is also anticipated in 2008, and has a peak throughput rate of 75 Mbps in 20 MHz channels. Clearly, CDG and Qualcomm are positioning EV-DO Rev B as an alternative to Mobile WiMAX.

While EV-DO Rev B represents an innovative extension of current CDMA technology, I do find it disturbing that CDG is pushing the 73.5 Mbps value, as this rate is an absolute peak theoretical value that represents the full capacity of 20 MHz of spectrum delivered to a single user under the most ideal conditions. Not once anywhere in the press release does CDG even mention average/typical throughput rates. Every value quoted is the highest theoretical value available.

We have been living with peak speed claims of 2.4 Mbps peak for EV-DO for a couple of years now. Having been involved in numerous field tests of this technology for a number of different organizations, I can assure you nobody has ever experienced 2.4 Mbps in the field. 1 Mbps is closer to the truth for peak measured throughput, and operators like Verizon are quoting 400 to 700 kbps typical throughput, about 20% of the quoted peak rate. I guess we better get used to the new hype value of 73.5 Mbps, as I’m sure it will be around with us for some years to come.

Mobile Commerce Rebound

Mobility Loop, Nov 18, 2005

Several years ago, there was a lot of buzz about mobile commerce, but then nothing substantive happened, and things went quiet. Now the buzz is back, and this time is armed with new methods. By mobile commerce, we mean being able to purchase items and execute transactions using your mobile phone. Users already do this, in a limited fashion, when they purchase ring tones using their handsets,

or if they are a Sprint customer, when they buy a song from Sprint's music service for $2.50. However, these are closed systems.

Where things get more interesting, and much more complicated, is with third-party financial systems and third-party products/services. The idea is that your phone becomes your wallet, so when you check out at a store, you could just pass your phone over the payment scanner and instantly execute the payment. This is the approach that NTT DoCoMo (Japan's largest cellular operator) just announced last week, with service to become available December 1 this year. The service will require new mobile phones equipped with wallet functions, and using Near Field Communications, the same technology as RFID, which allows devices in close proximity of 10 cm to communicate with each other. This helps address security and usability issues. In contrast, microbrowser-based approaches over cellular connections typically involve too many steps, and Bluetooth is not well suited for adhoc short-term connections.

There are no shortages of related efforts, nor failures, in this area. In Canada this week, Bell Mobility, Rogers Wireless and Telus Mobility this week announced a joint venture called Wireless Payment Services to provide a platform for wireless payments. Meanwhile, a European initiative called Simpay recently disbanded its initiative to provide a standardized platform for European mobile commerce. I expect massive confusion and conflicting approaches for quite some time, but the convenience of faster transactions, coupled with other benefits such as electronic coupons and expense tracking, will undoubtedly make mobile commerce a reality one day, with plenty of involvement by banks, credit card companies, operators and merchants. But the complexity of coordinating so many different parties, not to mention the difficulty of changing people's buying behavior, will make this a long slog.

Confusion About 4G

Network Computing Mobile Observer, Nov 9, 2005

Even though the rollout of 3G networks such as UMTS, HSDPA and EV-DO is only now gaining momentum and service is not expected to be widely available until the end of next year, there is already a lot of discussion about 4G as well as a considerable amount of confusion about what 4G really is. This is an area I've been looking at closely over the last year, including publishing a new public report on Oct. 24, 2005, titled "Hard Numbers and Experts' Insights on Migration to Evolved 3G and 4G Wireless Technology," presenting at the WiMAX World Conference on Oct. 26, 2005, on "Evolution of Mobile Broadband Wireless" and publishing a paper via 3G Americas in September 2005 called "Data Capabilities: GPRS to HSDPA and Beyond." These projects, among others, have given me the opportunity to collect facts and viewpoints on the current situation going from 3G to 4G wireless.

So it was with interest that I read on Oct. 19, 2004, that T-Mobile in Europe was deploying a new wireless service in the Czech Republic called "Internet 4G." Since everything else I had learned about 4G pointed to deployment after 2010, I was quite excited to think that somebody had leapt into the future, grabbed some great wireless technology to deploy today, and obsoleted global 3G deployments. But

that was not the case. In fact, T-Mobile's deployment is based on UMTS-TDD (Universal Mobile Telecommunications Network-Time Division Duplex), an official 3G technology specified by 3GPP (Third Generation Partnership Project). The vendor is IP Wireless, one of the world's leading UMTS-TDD equipment manufacturers.

When I questioned IP Wireless about the use of the term "4G" for this service, the company said it was the operator's decision to convey the sense that the service would provide greater capability than existing 3G networks such as UMTS. In fact, UMTS-TDD as implemented by IP Wireless is very similar to the HSDPA (High Speed Downlink Packet Data) version of UMTS operators like Cingular are beginning to roll out. The key difference is that most cellular operators are deploying UMTS/HSDPA in paired spectrum, with separate transmit and receive frequencies, whereas UMTS- TDD operates in a time division duplex mode. Since UMTS-TDD and HSDPA significantly outperform original data capabilities of UMTS, which is an official 3G technology, there is a desire among marketing types to differentiate the new services. But referring them to 4G is inaccurate. Here's the current roadmap to 4G.

First, operators are busy deploying HSDPA as well as CDMA2000 EV-DO (CDMA2000 Evolution Data Optimized). These highly capable services will offer throughput rates in the 500 to 1 Mbps range and should be widely available in most major U.S. metropolitan areas by the end of 2006. Then there are a series of enhancements planned for these technologies, including higher speed uplinks, advanced antenna architectures to boost speeds and improve network capacity and, eventually, VoIP in conjunction with quality of service control. Those items will happen over the next two to three years.

After that, 3GPP (representing GSM/UMTS companies) has an effort called LTE (long term evolution), which has goals of supporting flexible channelization (1.25 MHz to 20 MHz), peak throughput rates of 100 Mbps on the downlink and 50 Mbps on the uplink, extremely low latency and equipment ready for deployment by 2008-- a somewhat ambitious time frame. It is likely, but not definite, that LTE will use an OFDM radio interface.

3GPP2 (representing CDMA2000 companies) also has an aggressive evolution path, including EV-DO Revision B, which combines up to 15 radio carriers (each 1.25-Mhz wide) for a peak downlink speed of 46.5 Mbps), and a longer term project called 3GPP2 Evolution. The intent of both 3GPP and 3GPP2 efforts is to counter the competitive threat of Mobile WiMAX, which is much further along in terms of standardization but has little mobile operator commitment so far.

It is only after we get through this evolved set of technologies, which some have referred to as "Super 3G," that we start approaching 4G. There is no official 4G standardization yet, but early "industry consensus" indicates 1-Gbps peak rates, typical throughputs of 100 Mbps, network agility and best-of-breed radio techniques. The current 3GPP timeline for 4G deployment is around 2012, although I consider that a best case given all the interim steps.

Where does Mobile WiMAX fit in? Some have mused about WiMAX being a 4G technology, but that also is inaccurate. Mobile WiMAX does, however, embody many of the capabilities planned for evolved 3G systems. Hence, from a technology point, it is on an unofficial "roadmap" moving toward 4G.

Bottom line, other than companies co-opting the term "4G" for marketing purposes, we're a long way from 4G.

WiMAX World

Mobility Loop, Nov 3, 2005

This is my first piece for Mobility Loop, so it's a pleasure to write for you. I've been consulting in the wireless arena for ten years, and it's been an exhilarating non-stop ride of forever-evolving technologies. In the wide area space in the last decade, we've gone all the way from DataTAC and Mobitex to CDMA2000 Evolved Data Optimized (EV-DO) and High Speed Downlink Packet Access (HSDPA), with at least half a dozen intervening wide area technologies, some of which are still around. Most recently, WiMAX has received a lot of buzz as the possible "next big thing", so it was with interest that I participated in last week's WiMAX World conference in Boston. There was no shortage of interest in the topic, with at least a couple of thousand attendees and presentation rooms spilling over with people.

The solid news was the emerging availability of WiMAX-certified products for fixed wireless applications, and operators like AT&T (just acquired by SBC) describing their WiMAX trials and motivation for considering the technology, including bypassing the $8 billion a year that AT&T pays local phone companies for access to their customers. Also solid was the imminent completion of the IEEE 802.16e specification that adds mobility capability, with major vendors planning chips for the 2007 time frame. Less solid was what type of operators would deploy mobile WiMAX networks in what spectrum, to what types of users. However, one emerging view I discerned was on "personal broadband" data-oriented services in a three-tier network configuration, providing greater capacity than 3G networks though with smaller coverage areas, however exceeding Wi-Fi hotspot coverage areas. This makes sense, as Wi-Fi is somewhat problematic over wider continuous coverage areas.

One mesh Wi-Fi network vendor provided an interesting vision of mesh Wi-Fi networks becoming hybrid Wi-Fi/WiMAX networks, with WiMAX providing both access and backhaul, and access points deployed in the same locations as the mesh Wi-Fi network. Meanwhile, another vendor described a completely different model with WiMAX as a data overlay for cellular networks, deployed on existing cellular towers. There'll be no shortage of ideas thrown around in the next couple of years as the industry tries to figure out how to leverage this new wireless technology.

Wireless Data Pricing--One Hundred Times Less Expensive and Still Not Satisfied

Network Computing Mobile Observer, Oct 19, 2005

Verizon Wireless just recently reduced its unlimited data plan from $80/month to $60/month. Cingular quickly followed, and Sprint is likely to do the same (at the time I wrote this column, Sprint's Web page still shows $80/month for an unlimited data plan). This is a great step forward for many business users, but the service is still too expensive for the majority. It makes me wonder whether general-purpose IP-based cellular data can ever be priced at a level where such services will really take off.

Ten years ago, usage-based pricing for CDPD (Cellular Digital Packet Data) was 10 cents per kilobyte, which translates to a hefty $100 per megabyte. Today, usage-based plans run around $1/MB. So for usage plans, data pricing is one percent of what it was 10 years ago. That represents tremendous progress.

Yet, if I look at my own data usage, this is still too expensive for many scenarios. I don't travel enough to justify an unlimited plan with a PC Card modem. Instead, I have a usage-based plan with an EDGE (Enhanced Data Rates for GSM Evolution) setup that I use with a Treo 650 smartphone. What I find is that my smartphone (mostly e-mail and limited Web usage) nicely sips data, and my daily usage is in the hundred kilobyte range. But when I make the phone a modem for my laptop via Bluetooth, the laptop guzzles data at a scary rate.

I can, however, control the consumption by doing things like disabling Windows auto update, disabling auto virus signature file updates, and downloading e-mail headers and selectively viewing e-mail. By doing these things, my usage goes from about 10 MB/hour in normal office mode to about 1 MB/day in econo-mobile mode. I force myself through these contortions to better understand application data demands, not because I expect anybody else to do the same. In fact, it is this huge unpredictability about data consumption that makes unlimited data plans popular with companies. Especially when data overage is charged at more than 200 percent of the original rate (and international usage at 1,000 percent!).

So let's look at an uninhibited business laptop user making heavy use of wireless data. Assuming 50 MB/day and 20 days usage a month, this translates to a whopping 1 GB/month, equivalent in network usage to about 10,000 minutes of voice. That's a lot of talking, effectively 10 times the minutes of a typical voice plan. In other words, the cellular operator has to deliver 10 times its network capacity for data usage as it does for voice usage for about the same monthly revenue! If that same user was downloading a lot of multimedia content, data consumption could go up even further.

As detailed in a series of technical reports that I've published, the data capacity of cellular networks is quite limited. It's simply not feasible for operators to support a large percentage of their subscribers consuming gigabytes of data each month. Hence their reluctance to lower their unlimited data plans. The exception is T-Mobile, which offers unlimited data at $30/month. However, T-Mobile does not emphasize laptop or business connectivity, and makes most of its data revenue from services such as RIM Blackberry, Sidekick, SMS and MMS. The challenge for cellular operators is to find a tricky balance between pricing that encourages use but does not saturate the network. It is easy to criticize operators'

data pricing being too high, which it is for mass adoption, but being too aggressive about pricing could result in substandard service that would satisfy people even less.

So the emerging cellular operator strategy is to charge for value, not bytes. SMS is a terrific example. At 10 cents per message for a pure usage-based plan, the operator is collecting about $1,000 per MB. Hosted e-mail solutions are another good example of selling value instead of packets. Look to operators to use their data networks for all kinds of other value-added services rather than just IP connectivity, as it looks like there might be little real money to be made in just selling fat dumb pipes to mobile users.

Embedded Wireless WAN

Network Computing Mobile Observer, Sep 28, 2005

One important enabler for the WLAN industry has been the embedded WLAN capability for notebook computers, which is now included in most of the units sold. Intel's Centrino effort has both provided marketing dollars and made it easy for notebook computer vendors to include WLAN capability. Can the same success be replicated with wireless wide-area networks, whether they are based on EDGE, EV-DO, UMTS/HSDPA (High Speed Downlink Packet Access) or, eventually, WiMAX? Will such developments also accelerate the adoption of wireless data? And what are the issues involved in making embedded wireless WAN a reality? These questions were addressed in a September 13, 2005, symposium by the Portable Computer and Communications Association, an organization that works on wireless data interoperability issues and of which I act as chair.

There are a number of reasons for wanting wireless WAN embedded. One is ease of use, since the card comes preconfigured for use and doesn't require an installation procedure. Another is improved performance. If the notebook computer is properly designed for the wireless WAN card, it will include internal antennas high in the lid of the notebook computer. It may even include multiple antennas to accommodate forthcoming performance-increasing techniques such as mobile receive diversity and MIMO (multiple-input multiple- output). The notebook computer will also provide the proper thermal, electrical and electromagnetic environments to ensure optimum operation of the radio circuits.

Many companies are working to make embedded wireless WAN a reality, among them chip vendors, modem vendors, notebook manufacturers and cellular operators. So far, only a few solutions are available, most notably the Sony Vaio with EDGE capability using Cingular's network. However, a number of recent announcements point to developing momentum. Dell recently announced it would offer an EV-DO PC Card as well as provide service activation for Verizon Wireless. For the moment, this is a separate PC Card, but this will likely migrate to an embedded format next year. And last week, Dell announced that beginning in the first quarter of next year it will offer notebook computers with embedded HSDPA capability using Cingular's network. HP, Lenovo (IBM ThinkPad) and Panasonic have also indicated plans for supporting embedded wireless WANs. For embedded installations, the emerging form factor is PCI Express Mini Card. Vendors will be shipping modems in this form factor in 2006.

The challenge is that the market is somewhat fragmented with multiple wireless technologies. For embedded wireless WAN to work the most effectively, notebook computers must be designed essentially from the ground up to support the addition of wireless modules, including internal antennas, the appropriate electronic designs and board layouts to minimize interference from the computer, and appropriate heat dissipation. Complicating things further is that different operators have different requirements from their vendors. This means the EV-DO modem for Verizon is not necessarily identical to the EV-DO modem for Sprint. However, issues like this are largely manageable and solutions are facilitated by the consolidation of wireless operators.

The industry is also addressing how to manage wireless transmitters in devices for usage in environments such as airplanes and other areas where transmitters must be disabled. For instance, the Consumer Electronics Association has published a document entitled "Recommended Practice--Status Indicator for and Control of Transmitters in Portable Electronic Devices (PEDs)."

With sales of wireless WAN cards much lower than wireless LAN cards, it is doubtful that wireless WAN capability will become a standard item on notebook computers for some time. Incremental hardware costs for wireless WAN are currently much higher than for wireless LAN, where costs are in the $10 range for IEEE 802.11g modules. Wireless WAN costs today are at least an order of magnitude higher. Over time, this will change. Intel is pushing to make mobile WiMAX a standard feature on future notebook computers, and competing EV- DO and HSDPA forces will attempt to do the same. But we are at least a couple of years from this scenario.

Nevertheless, the capability doesn't have to be standard for the embedded capability to provide significant benefits. Making wireless WAN a checkbox item when companies order notebook computers increases awareness of the wireless data services, simplifies purchasing and provides better performing solutions. Bottom line: Embedded wireless WAN won't revolutionize the industry, but it is one important step toward the greater adoption of wireless data.

Redefining the Endpoint--Wireless Broadband Routers

Network Computing Mobile Observer, Sep 7, 2005

Things used to be simple with wide-area wireless networks. The network endpoints were mobile telephones or modems. The modems were inside PDAs or laptops. The endpoints were basically you and me--people. But that's changing, with companies like Netgear, which last week announced a wireless broadband router for Flash OFDM. Netgear's router has a wide-area connection using Flarion's (now owned by Qualcomm) Flash OFDM technology and local connections using Wi-Fi and Ethernet. Okay, so there are no Flash OFDM networks in the United States. But Cingular and Sprint have offers in this area, and Verizon is not far behind. The endpoint is no longer a person; it's a network. And just as we were going to press, we saw D-Link's August 30 announcement of its Wireless 3G Mobile Router.

The general idea is to terminate the wireless WAN connection at a device that routes to a network. I got a nice taste of this recently at an industry association meeting where an enterprising network engineer had a UMTS card in his laptop and distributed his Internet connection over Wi-Fi to willing recipients in the meeting. About a dozen of us enjoyed this ad hoc service, and it worked just fine. In this case, Microsoft Windows in the laptop was doing the routing, with each of us using a separate IP address provisioned from the laptop.

So should we all start using our UMTS (and forthcoming HSDPA) or EV-DO connections as an alternative to T1s or DSL connections and share that connection with multiple endpoints? However tempting that may sound for some scenarios, it's not the best application of the router technology. Wireless broadband speeds are on average lower than wired connections, latency is higher and, most importantly, monthly costs are higher. Furthermore, the unlimited usage service plans are likely to prohibit this type of usage. The application that does make a lot of sense, however, is backup connectivity, particularly for mission-critical situations. Rather than having a redundant wireline connection, you can use the wide-area wireless connection for backup.
For applications such as metering and retail point-of-sale that don't transmit much data, the wireless connection might even be the primary application.

Cingular supports this type of connectivity with its Wireless WAN Connectivity service, which is accompanied by a router from Digi called Digi Connect WAN GSM. This router is a compact device that supports GPRS and EDGE. Sprint's offer in this area is its PCS Data Link service, which originally supports 1XRTT, and will be an obvious candidate for Sprint's forthcoming EV-DO service. Verizon does not yet have a formal plan. A number of routers support EV-DO; for example, the Kyocera KR1 router with Wi-Fi and the TopGlobal routers. Junxion makes a router that allows you to insert PC Card format modem cards and supposedly support EDGE, 1xRTT and EV-DO. The D-Link router also accepts modem cards for all the major 3G networks. Omniwav has an EV-DO router designed to be installed in vehicles. Every major router vendor is likely to follow suit. This market area is in its early days, but clearly coming alive.

If you're thinking about how to take advantage of wireless WAN network-to-network connectivity, here are some considerations. First and foremost, make sure the service plans allow this kind of usage. Then make sure the wireless WAN connection has the horsepower for your application. Realize that wireless speeds vary according to signal quality and network loading, so test the application in both best-case and worst-case scenarios. Then investigate pricing. Even though operators have unlimited usage plans, these are designed for end users and are not intended for transmitting large numbers of gigabytes of data every month.

As impressive as the capabilities are of new wireless networking technologies, I have previously commented on their relatively low capacity, meaning it doesn't take very many endpoints communicating continuously to swamp the capacity of the local cell site. Operators are very sensitive to this, and it is reflected in their pricing plans. For instance, Cingular offers usage based plans from 5 MB for $19.99 per month to 50 MB for $49.99 per month and also allows multiple devices to be pooled together to determine total monthly usage, which can be useful as it can be hard to predict the amount

of traffic at any one node. Sprint's pricing is at a similar level, with 40 MB costing $40 per month. You'll also need to estimate how much traffic may go over your backup connection so you don't get nailed with overage charges.

You'll still have to do your network planning, such as designing routing approaches. The Digi product, for example, supports GRE encapsulation, which allows the remote network to have IP addresses assigned via the home network.

With the higher networking speeds of technologies such as EV-DO, HSDPA, Flash OFDM and WiMAX, the wireless broadband router market makes a lot of sense, and it is likely to become one more tool in your networking toolkit.

High Speed Downlink Packet Access (HSDPA)

Network Computing Mobile Observer, Aug 17, 2005

First, an important correction: In my last column on wireless data uptake, I reported 500,000 EV-DO customers for Verizon--the number that has circulated in various stories. Mobile Pipeline Editor Dave Haskin questioned this number and obtained a clarification from Verizon that the number refers to data subscribers for both 1xRTT service and EV-DO. Most likely then, the majority of the customers are using the earlier 1xRTT service. Now, on to HSDPA (High Speed Downlink Packet Access).

Vodafone, one of the world's largest cellular operators, and Nokia just completed testing of HSDPA, a new version of WCDMA (Wideband CDMA), in Italy using Nokia infrastructure. The two companies reported throughput rates of 1.5 Mbps. Various other companies have demonstrated throughput speeds of over 1 Mbps for this technology, and peak speed claims of over 14 Mbps are also common. But what speeds can you realistically expect and what services will be available when? Also, how does this technology line up with the competition, including EV-DO, Flash OFDM and WiMAX?

First, let's look at what the technology is. HSDPA is an enhancement to WCDMA, a technology that is also referred to as UMTS (Universal Mobile Telecommunications System), the 3G path chosen by most GSM operators around the world. Today, there are already some 75 UMTS networks in operation around the world. In the United States, Cingular (via its purchase of AT&T Wireless) has UMTS available in six cities.
However, the company is planning an aggressive deployment of HSDPA, with 15 to 20 cities planned by the end of 2005 and most major metropolitan areas by the end of 2006. The six cities with current UMTS service will be upgraded to HSDPA as well. Globally, many UMTS operators are planning on the HSDPA upgrade, and operators that have not deployed UMTS yet are likely to go directly to HSDPA.

HSDPA is important for operators because current UMTS technology, which based on 3GPP WCDMA Release 99 specifications, is extremely efficient for voice service, but it is not optimized for data services. The HSDPA upgrade, based on 3GPP Release 5 specifications, keeps the same voice mechanisms but

adds highly efficient data capability. It does this through radio mechanisms such as higher order modulation (16QAM in addition to QPSK), improved error correction, dynamic adaptation of modulation and coding based on radio conditions, and a tighter link between mobile and base station. It also adds a packet scheduling approach that favors downlink transmissions to users with the best radio conditions, thus allowing them to receive data at higher throughput rates. Since conditions vary by users over time, this results in what is called user diversity. All these approaches working in combination result in at least a doubling of spectral efficiency, and with later improvements on the roadmap, a further doubling. Bottom line: The network can support many more users at a higher speed.

HSDPA is not the first technology to invent these approaches. You see them in most new wireless technologies, including CDMA2000 EV-DO and WiMAX, as well as in Wi-Fi. As for speeds, that's where things get a bit tricky. The network does indeed support a peak rate of 14 Mbps, but this is what you would get with the full capacity of the cell sector at the highest modulation level and no error correction. In other words, you'll never actually experience this speed, as you'll never experience the peak rate of 2.4 Mbps quoted for EV-DO. Initial devices are likely to have peak rates of 1.8 Mbps or 3.6 Mbps, depending on their designs. Subsequent devices will be faster. As for typical speeds you can expect, it will depend on the number of users active in the network, but average speeds in the 500 kbps to 1 Mbps range should be achievable in relatively lightly loaded networks.
However, speeds may go down a bit if people flock to the technology, as is the case for all the 3G technologies. Of course, with the somewhat slow adoption of cellular data as discussed in my last column, this is a problem that operators would love to have, and one they can manage through additional cell sites and new spectrum allocations. HSDPA latency goals are also aggressive, and in initial networks, latency will likely be in the 100 to 200 msec range.

As for how this lines up with EV-DO that Sprint and Verizon are deploying, HSDPA will likely meet or exceed EV-DO performance. HSDPA also allows simultaneous voice and data on the same device. However, EV-DO operators will have broader coverage in the near term. Mobile WiMAX specifications point to somewhat higher performance, but there are a series of further enhancements planned for HSDPA.

In my view, HSDPA is a great thing for the wireless industry for a number of reasons. First, it will offer a highly compelling broadband wireless service over very large coverage areas. Second, it will compete head to head with EV-DO, resulting hopefully in competitively priced service plans. Third, it will keep the WiMAX people on their toes, forcing them to deliver the best possible technology; otherwise, it won't be competitive. The reciprocal case is also present, as WiMAX has accelerated 3GPP (the organization that specifies GSM-UMTS
technology) efforts to start designing what comes after HSDPA, currently called 3GPP Long Term Evolution.

From a usage point of view, even by the end of 2006, HSDPA won't be everywhere. Coverage is likely to emphasize urban areas, and in less densely populated areas, the fallback will be to EDGE for quite some time. This means there's no reason to wait for HSDPA, as you should probably plan for most applications

you deploy to work over both HSDPA and EDGE. The same is true for EV-DO, where the fallback is to 1xRTT.

As HSDPA gets closer, you're likely to see lots of hype about the technology. But if you keep in mind its true capabilities, there's no reason to not start incorporating HSDPA into your wireless planning.

The Uneven Uptake of Wireless Data

Network Computing Mobile Observer, Jul 21, 2005

Why is Wi-Fi surging ahead and why is cellular data still struggling, even against the backdrop of global 3G network deployments? The simple answer is that Wi-Fi is fast, cheap and easy while cellular data remains slow, expensive and complicated. That's beginning to change, but slowly. First some evidence to support the claim. A new study from Research and Markets projects that WLAN equipment and services will grow at an annual rate of 30 percent, to a global figure of $5 billion for the year by the end of 2006. In other words, Wi-Fi is continuing to grow at a healthy clip. Meanwhile, IDC in May reported that wireless data constituted only 5.8 percent of wireless revenues in the United States during the fourth quarter of 2004. This 5.8 percent actually adds up to $1.6 billion in revenues for the quarter, which is nothing to sneeze at, but it represents limited usage, especially when you consider that some half of this figure is revenue from text messaging. IDC expects this figure to grow to 20 percent by 2009. Let's try and understand the dynamics of this slow growth.

Like any successful technology, wireless technology must provide value without exacting too high a cost. That is why Wi-Fi is so successful and popular. It allows me to write this column in the comfort of my living room while providing all the broadband speed I need at a cost that is laughably low. And that is why cellular voice service has grown to 1.8 billion subscribers around the world. It offers the convenience of telephony from anywhere, at any time, at a cost that consists of the monthly fee--affordable for most people--and, with degraded voice quality relative to wireline, but at a quality level that most people accept.

Now compare this with cellular data. Let's start with the laptop scenario. There is definite value in having connectivity from anywhere, and I personally can't imagine giving up the convenience of my nationwide wireless data service. But today, the costs are still high. First there is the obvious monthly charge of $80 for EDGE or EV-DO for unlimited usage. This eliminates quite a few potential users right off the bat. But there are plenty of other costs. Let me list some. With non-3G technologies, the networks are still too slow for many applications. Chances are that your laptop is also configured for automatic Windows and virus updates, meaning your system may be downloading megabytes of information and you won't even realize it. This slows down the performance of other applications and drives up bills on usage-based plans. That's not the network's fault, but it's still a usage cost. Yes, there are workarounds, such as making sure laptops are configured exactly right for wireless networking, using wireless middleware, mobile-optimized versions of applications and so forth. But dealing with all these represents a large integration cost compared to just "using" your existing networking applications.

Meanwhile, configuring a phone to work as a modem via Bluetooth can be devilishly complicated. The alternative is a separate account and a data card.

Continuing on, Microsoft Windows XP has integrated support for Wi-Fi that provides a consistent user experience regardless of Wi-Fi hardware vendor. However, there is no comparable support for cellular data. And in the next version of Windows, Windows Vista, you can look forward to further improved Wi-Fi support and still no native support for cellular data. This doesn't mean you can't use cellular data with Windows, it just means it isn't as easy as it should be, for both end users and IT managers, and it requires additional utilities you have to run to configure and manage your connections.

In the smartphone scenario, including wireless PDAs, devices have impressive computing capability. But the number of business applications targeted for these devices is still very small, with wireless e-mail and calendar management seeing the greatest usage. Making other enterprise information available often involves custom development. More cost. And if you have a smartphone, chances are that the operator has neutered it so you can't use it as a modem for your laptop anyway. Why? Because of potential abuse by some users with lower priced smartphone plans consuming too much data using their laptops. This is the single biggest reason I don't use a smartphone, much as I would like to. Today's smartphones also often involve separate management and security architectures from their laptop cousins--another cost.

So is there any hope? I believe so. To begin with, new data technologies such as EV-DO and HSDPA are beginning to overcome the performance barrier. This means that many existing applications work at an acceptable level--though, ironically, only if the networks don't become too popular, as network data capacity is still relatively limited. Despite the high monthly cost, Reuters reports that Verizon now has 500,000 subscribers for EV-DO. Tiny compared to its 40 or so million voice subscribers, as of June, but not bad for how new the service is. Pricing is also improving, albeit slowly. For instance, Sprint's just released EV-DO pricing is $80 unlimited, but it also offers a $40 plan for 40 megabytes that is capped at $90 if users exceed their quota, a creative pricing compromise. I was briefed last week on Sprint's new program, called Sprint Mobile Business Assessments, where Sprint works closely with a company to help it develop comprehensive mobility solutions. This is typical of how all the operators are providing assistance to enterprises to help overcome integration challenges.

Smartphone-oriented applications such as wireless e-mail are also getting good traction. They provide an effective e-mail experience at a reasonable monthly fee and are easy to install, configure and manage. Other enterprise applications for smartphones are starting to see the light of day. So piece by small piece, the cellular data picture is getting brighter. But for somebody such as myself, who has been involved in wide-area wireless data projects for
14 years now, it's sure happening slowly.

Where is WiMAX?

Network Computing Mobile Observer, July 6, 2005

WiMAX deserves an award for having the highest visibility relative to actual availability of any wireless technology ever. More IT managers have heard of WiMAX, which is not yet available in any form, than some of the 3G technologies that are now available in many areas. And during June, WiMAX enjoyed additional press from a long series of announcements involving both service providers and equipment vendors. Among them, AT&T announced it would start a field trial of the fixed version of WiMAX in Atlanta this fall; the company is planning to evaluate the technology, the applications that are feasible and the business case. Qwest Communications indicated that it is currently evaluating fixed WiMAX technology and is planning a trial of residential and business users, supposedly in a location where DSL is not readily available. On the equipment side, RedLine Communications announced WiMAX products for the 3.5-GHz licensed band available for broadband wireless service in Europe and Asia. Equipment vendor Aperto Networks said it was doing interoperability testing with WiMAX chips from Intel and Fujitsu.

As for mobile WiMAX announcements in June, Nokia agreed to work with Intel to help complete the mobile WiMAX standard (IEEE 802.16e) and to eventually develop mobile equipment as well as infrastructure equipment to support the mobile WiMAX standard. This is notable because, until now, Nokia has had a fairly ambivalent stance toward WiMAX. Intel and ArrayComm, a leader in smart antenna technology, announced a collaboration to incorporate smart antenna capabilities into the IEEE 802.16 standard, with Intel planning to support ArrayComm's techniques in future IEEE 802.16e chipsets. Though few networks today employ smart antenna technology, it is viewed as an effective means of improving the performance and capacity of wireless networks. Navini Networks announced a line of products, including client and base station equipment, based on IEEE 802.16e. The client device will be available by the end of the year, with a PC Card modem and base station equipment following in 2006. Finally, Sprint entered into an agreement with Motorola to jointly test Motorola's equipment in its 2.5-GHz spectrum band. In May, Sprint had said it would work with Intel to develop mobile WiMAX capabilities and use them to deliver consumer-oriented multimedia services.

Additionally, we have the general development of a company doing anything with broadband wireless now calling its equipment "Pre WiMAX," lending WiMAX legitimacy by suggesting that it's not really broadband wireless if it's not WiMAX. To be fair, most of the companies using the "Pre WiMAX" term do intend to deliver WiMAX products.

So where is WiMAX really? WiMAX actually is in two very different places. The fixed version, based on IEEE 802.16-2004 (sometimes referred to as IEEE 802.16d), is materializing, with chipsets now available, vendors planning to do initial interoperability testing during the second half of this year and products becoming generally available during 2006. This is what the Qwest and AT&T trials will be using, with the most likely applications being DSL-type services in areas underserved by DSL and local telephone bypass, which will be a little trickier because it will probably require VoIP. Smaller ISPs may also jump onto this bandwagon, but since most don't have licensed spectrum, these companies will likely do this on a niche basis using the 5.8-GHz unlicensed band. Products in this band will also be available for private deployments. This is all real, and if the technology works as advertised, we may start to see real

commercial service from operators beginning in 2006. Initial products are likely to be line-of-sight, followed by non-line-of-sight options.

On the mobile side, things are fuzzier, as WiMAX is barely on the horizon. First, work on the specification continues, with hopes that it will be done by the end of this year. There is a lot of pressure to complete the specification, as many vendors view the mobile version as having much greater market potential than the fixed version. However, even in the best-case scenario, the earliest networks could become available is the end of 2007, with 2008 much more likely. One trend here is that many of the 3G vendors, especially on the infrastructure side, are becoming involved and planning to offer both 3G and mobile WiMAX products, thus hedging their bets. Another trend is that vendors are taking a more conciliatory approach with respect to 3G. Early views were that mobile WiMAX would outperform 3G, but now vendors such as Nokia are saying that WiMAX will be complementary with 3G, allowing, for instance, additional broadband capacity in hot zones for demanding multimedia-type applications. What will actually happen is far from clear. The cellular community is already working on specifying evolved 3G systems and experimenting with 4G concepts based on many of the same principles as WiMAX. Watching to see whether WiMAX co-exists, usurps or is made redundant by these developments will be a fascinating spectacle during the second half of this decade.

Wireless E-Mail for the Masses

Network Computing Mobile Observer, Jun 15, 2005

Wireless e-mail will eventually reach penetration levels close to that of wireless voice, with 80 percent of mobile phones being used for wireless e-mail. That provocative forecast was made to me during an interview with SEVEN as part of a series of interviews we conducted for a forthcoming article in Network Computing about wireless e-mail. I don't intend to give away our conclusions in that article prematurely, but I do want to examine more closely the aspect of how far wireless e-mail might go, as we didn't have space for that in the article.

At first, I was skeptical of the 80-percent penetration figure. But I've given it considerable thought, and numerous recent announcements shed some insight into the way the market is evolving and provide evidence that this market has tremendous upside. On June 6, 2005, Microsoft unveiled its Messaging and Security Feature Pack for Exchange Server
2003 and corresponding support in Windows Mobile 5.0 for direct wireless e-mail push technology. On May 31, 2005, Good Technology announced a deal with Cingular to provide its wireless e-mail gateway technology as an operator-hosted solution that does not require companies to purchase a behind-the-firewall gateway from Good. And then on June 8, 2005, what really got my attention was Sprint and Yahoo announcing a push e-mail service for Sprint subscribers that supports Yahoo mail. (Yahoo has some 60 million e-mail subscribers.) The monthly cost for the service is only $2.99.

I have also recently experimented with Yahoo's SMS notification service for e-mail, a free service that allows you to configure Yahoo mail to apply a set of filtering rules to your incoming e-mail and to send

you text message notification of filtered messages. Don't use Yahoo mail? No problem, as long as you have a POP3-accessible mail account somewhere. Just get a free account and have Yahoo retrieve your mail using POP3, and then you'll receive notifications. AOL and MSN have ever-increasing mobile options for their subscribers as well. These market developments point to an increasing number of service-side options for both business users and consumers that will make e-mail readily accessible and convenient using mobile devices.

The other side of the equation is developments on the device side. Whereas smartphones and PDAs are typically what you associate with wireless e-mail, the fact is that an increasing number of ordinary phones are also wireless e-mail capable. I use a Nokia 6230, a nice capable compact phone based on Nokia Series 40 platform software. This is not a 'smartphone,' yet it is smart enough that I currently have three options for accessing e-mail. One is to use the WAP browser in conjunction with a Visto-powered operator service for viewing my e-mail. I can use an e-mail client built into the phone to retrieve POP3-based e-mail. I also installed a more fully featured Java-based e-mail client, EmailViewer from Reqwireless, which has some additional features beyond the built-in version.

Each e-mail approach has its limitations compared to using a larger smartphone with micro-keyboard. But for many scenarios when I'm already traveling with a laptop, this is all I really need to monitor my inbox for urgent items. The point is that, already, practically every mobile phone has some means of wireless e-mail. With increasing memory size, Java or BREW capability, in particular, and better screens, wireless e-mail will be a button press away for most users. Factor in also that smartphones costing $400 or more today will likely be $200 in a year or two.

So, we have the right server-side developments, we have the right device-side developments and, through the Sprint/SEVEN service and likely promotions, we'll have increasing user awareness. The question then is, who would want to use this service? My answer is, who wouldn't?
People who use e-mail on a regular basis and who also use a mobile phone are likely to find themselves in situations where browsing e-mail on their phones will be a convenient option. Real-time e-mail capability may not be essential to your life. But just as you may check your home voicemail while on vacation, you may also want to check your e-mail. Or maybe it is just something to occupy your time while waiting in line at the bank or for a movie to begin. Either way, all the pieces are falling into place for making wireless e-mail available to the masses.

Parallel Mobile Universes

Network Computing Mobile Observer, May 25, 2005

Any company wanting to broadly deploy mobile or wireless applications has to deal with two completely different sets of communications infrastructures and, in many cases, both are used to access the same corporate information. This has crystallized for me as I've been completing a number of projects, one being a white paper for an SSL VPN vendor on the benefits of using SSL VPNs to secure wireless communications and the other an article for Network Computing on wireless e-mail. Here's the crux of

the problem: The way we want to have information delivered to us on a handheld PDA or smartphone is fundamentally different from what we want on a notebook.

On a notebook, mobile workers are essentially replicating a desktop experience. They are using many of the same applications--e-mail, databases, client/server, Web browsing, group productivity and so forth--they use in their office. The usage model is usually a nomadic one, with the user turning on his or her computer in some location and then initiating the data transfer. With a handheld device, the primary application today is e-mail, followed by calendar synchronization and, increasingly, access to bits and pieces of corporate information. So far so good: These same applications also run on the notebook. What differs, however, is the usage model. What people want with e-mail and other mobilized information is the ability to push the information to the mobile device. Push allows for real-time communications; for example, an urgent e-mail or instructions telling a mobile worker what to do next. This ability not only has made RIM so successful with the BlackBerry, but it also has helped to develop an entire wireless e-mail industry with a variety of competitive providers.

As it turns out, this little matter of "push" changes everything. Most client/server applications, whether e-mail or Web browsing, are based on the client initiating the data transaction over the network. But with a solution like RIM's, you install a separate server behind your firewall that accesses e-mail and other services on behalf of the mobile device and then sends (pushes) new information to the mobile device, sometimes via a network operation center (in the case of RIM and Good Technology) and sometimes directly to the operator. For privacy, the server and mobile device use encrypted tunnels. Since the loss of a mobile device itself can be a security hazard, most of the mobile solutions allow an administrator to initiate special messages that do things like delete data on the mobile device, reset passwords or turn off the wireless capability. The net result is that these mobility solutions, designed initially for optimized push e-mail, now offer a fairly comprehensive security architecture. But how does this security architecture relate to the security architecture being used for notebook computers? It doesn't, and therein lie the parallel universes.

With notebook computers, the security architecture is based on a secure extension of general-purpose IP-based communications, and the approaches used involve IPSec VPNs, the increasingly popular SSL VPNs and mobile VPNs designed specifically to optimize wireless networking. Although you can, and people do, run VPNs on PDAs and even smartphones, this approach doesn't give you the push features that people find so desirable. It also encumbers sessions with long connection times and sensitivity to dropped tunnels because of the fickle radio environment. As a result, the VPN-based security approach for remote access operates completely independently of the communications infrastructure used for handheld devices.

This duality is unfortunate, as it places an additional burden on IT managers for supporting mobile workers. In an e-poll for my Network Computing article on wireless e-mail, respondents expressed a relatively strong preference for wireless e-mail solutions that use existing VPN systems for remote access to servers.

I have asked a number of experts in both the wireless e-mail and VPN communities whether they see any convergence on the horizon, and the answer has consistently been no. The reality is that the notebook computer is essentially a mobilized desktop computer, whereas the handheld device is an entirely new type of platform with unique communications and usage requirements. Today, the penetration of mobile voice is extremely high, but the penetration of mobile data is extremely low. Yet it seems inevitable to me that, over time, the percentage of people using their mobile devices, be they cell phones or wireless PDAs, to obtain important information in real time will become extremely high--well over 50 percent. Hence, the need for secure mobile communications architectures will only grow. Now if only we could learn to turn the darn things off now and then.

Beyond the Airlink- IP Multimedia Subsystem

Network Computing Mobile Observer, May 4, 2005

Wireless technologists like myself tend to get caught up in the various advances of radio technology, and for good reason. There are a lot of fascinating developments, and better radio performance translates to better user experiences and a greater range of available applications. But there are also developments at higher levels that may have even greater implications on the types of services you can expect from wireless networks over the next several years. One such development is IMS (IP Multimedia Subsystem), which is being specified by 3GPP (Third Generation Partnership Project), the group that maintains the GSM/UMTS family of specifications.

IMS is essentially a service platform for IP multimedia applications. This includes potential applications such as PoC (push-to-talk over cellular), VoIP, streaming video, video conferencing and so forth. For example, you're on a voice call and suddenly want to enable a video connection or transfer files. Or maybe you're on an interactive chat session and want to use the chat session to launch a voice call. Or you're browsing the Web and decide you need to speak to a customer-service representative. Maybe you're on a business trip and want to stream video to your handheld device from your home digital video recorder. Your teenage sons are having an interactive video chat, and then launch an interactive game while continuing their chat. Think of any creative blending of voice/data/video and you're thinking IMS. It's no wonder that IMS has been such a big buzz at recent wireless conferences.

IMS by itself does not provide all these applications. Rather it provides a framework of application servers, subscriber databases and gateways to make them possible. The exact services will depend on operators and application developers who make these applications available to operators. For the most part, these won't be applications that are enterprise-hosted. However, the operator services will be able to intelligently interact with Internet-based and enterprise-hosted services because IMS is based on standardized networking protocols.

The core networking protocol used within IMS is SIP (Session Initiation Protocol), which uses its companion protocol SDP (Session Description Protocol) to convey configuration information such as supported voice codecs. Other protocols include RTP (Real Time Protocol) and RTSP (Real Time

Streaming Protocol) for transporting actual sessions. Though specified by 3GPP, IMS is relatively independent of the radio access network and can, and likely will, be used by other radio access networks such as CDMA2000 and WiMAX. It can even be used with wireline networks such as DSL and cable. Operators are already trialing IMS, and PoC is likely to be one of the first common IMS-based applications. Operators looking to roll out VoIP over networks such as CDMA2000, EV-DO and HSDPA will also use IMS.

The fact is, nobody quite knows what the killer application for 3G will be. But operators see IMS as a strategic tool to enhance their revenues from 3G services. This will be especially important to avoid evolving into a commodity business selling just IP connectivity; otherwise, operators' fates may be similar to what happened to long-distance telephone carriers.

There is, however, a lot to be done to make IMS work. First up is realizing that IMS itself is a work in process. For example, it is spread out across two releases of 3G specifications. The first, in release 5, is for non-real-time or near-real-time services. This release corresponds to the same release that is bringing out HSDPA to the world, namely high-speed packet-data services for UMTS with initial deployments at the end of this year. Enhancements for real-time services will come in release 6 specifications, which correspond to increases in uplink speeds for UMTS. It may be 2006 or later before this version is available.

One big missing piece for IMS is QoS (quality of service), which will have to be managed both in the IMS infrastructure and in the radio-access network. You can't do VoIP or video conferencing over wireless using a best-efforts approach.
You have to have QoS. The various 3G technologies all have QoS mechanisms defined, but by and large operators have not yet turned on these features.

It's unlikely that operators will call their services IMS. But as you see innovative new services emerging beginning this year and next, it's likely that they will be based on IMS. If you have large-scale plans or deployments using SIP-based services, you might want to find out what your operator is planning in the area of IMS as there could be some great future synergies. While it may take the rest of the decade before we move to a complete IP-based telecom infrastructure, the foundation is being built today from pieces like IMS.

The Wireless Edge: 3 Sides to 3G.

Network Computing Mobile Observer, Apr 13, 2005

3G wireless technology has tremendous potential, but its biggest threat could be its own success. Its success, however, will lay the foundation for the next generation of wireless technology. I've been researching and writing about wireless data for over 10 years and have been through Mobitex, CDPD, GSM circuit-switched data, GPRS, EDGE, CDMA QuickNet Connect and CDMA 1xRTT. In U.S. markets, the excitement now is over 3G technologies such as UMTS/WCDMA/HSDPA (Universal Mobile

Telecommunications System/Wideband CDMA, High-Speed Downlink Packet Access) from Cingular and CDMA 2000 EV-DO (Evolved Data Optimized) from Verizon and Sprint. And even these technologies are not standing still; Verizon last week indicated it would start field trialing an upgrade to EV-DO called Revision A around the end of 2006.

There is every reason to be excited. With typical throughput rates of around 500 kbps and latency in the 200 to 300 msec range (and moving downward with HSDPA and EV-DO Revision A), many networking applications that were quite sluggish with previous networks now perform quite well. Even operating IPSec and SSL VPNs over 3G networks is quite feasible despite their networking overhead. With planned enhancements, we're looking at typical rates closer to 1 Mbps and latency approaching 100 msec. It's not quite the broadband experience of a Wi-Fi connection--for instance, full-featured Web pages can still take tens of seconds to download--but it's not that far off either. At a recent industry association meeting I attended, an enterprising colleague had a UMTS/WCDMA connection that he allowed others in the room to share via Wi-Fi. About a half dozen of us took advantage of this, and we all had reasonable performance for e-mail and Web browsing, and that was over just one UMTS connection.

Within a year or so, these 3G services will be available nationally, and with competition, there will be downward pressure on pricing currently hovering around $80 for unlimited use. That's the good news. The second part of the story, the bad news, is that despite the excellent performance, the current macrocellular architecture inherently has limited capacity. This has to do with how many users can access the available radio spectrum simultaneously for voice and data given large cell sizes.

In an analysis I completed recently on the spectral efficiency and capacities of different wireless technologies, it was quickly clear that broadband data demands could rapidly swamp these networks. While monitoring my own network usage, I observed that I can easily consume 10 MB of data per hour on the downlink through typical e-mail usage, exchanging documents with colleagues, accessing databases and doing research on the Web. Taking a highly evolved cellular technology such as UMTS/HSDPA, one cell site with one 5-MHz radio channel can support about 3,000 users who, each hour, consume 10 minutes of voice service and 100 KB of data (indicative of wireless e-mail).

But if those users are doing office applications and consuming the 10 MB of data I just quoted, then that figure quickly falls to about 500 users. These numbers of users assume a highly uniform demand distribution. Given that user data traffic is bursty, the actual number of users who can obtain satisfactory broadband service actually will be a much lower number, though I haven't seen an analysis yet that fully accounts for typical wireless data usage patterns. With a coverage radius of 1 kilometer, the limited network capacity translates to just a small percentage of the total population being able to obtain broadband service from such a network.

That's the conundrum. If the cellular operators ramp up broadband subscribers, they can quickly consume the network. This is one reason operators are not in any hurry to decrease their flat-rate pricing plans. Wi-Fi gets around this capacity problem by allowing the same spectrum to be reused over much smaller geographic areas.

This leads to the third part of the story: In my optimistic view, 3G data services for both business and consumer entertainment and multimedia applications will see steady growth and will start to generate significant revenues, thus providing the impetus for the next generation of wireless technology deployment. What is that exactly? It's not completely clear. Options include new 3G spectrum that the FCC plans to license in mid 2006, denser deployments of base stations and enhanced radio techniques such as antenna diversity, channel equalization, MIMO, multi-carrier CDMA and OFDM. Down the road, MIMO-OFDM is a promising candidate for maximizing spectral efficiency in wide radio channels. Another approach is closer integration between 3G and Wi-Fi networks, which would allow Wi-Fi to serve high user densities in places like office buildings and airports. It's likely that vendors and operators will use all of the above in creative combinations. Coupled with this next phase of wireless technology, you can expect a move to all IP networks with voice handled as VoIP.

Successful business models for this new wireless future remain elusive, yet the momentum is huge. And keep in mind that it will be several years before we get to this next stage. In the meantime, assuming you can afford them, there are few reasons not to start using 3G services. They are very good, getting better and have a very bright, though somewhat out-of-focus future.