

The Internet Protocol Journal

June 2008

Volume 11, Number 2

A Quarterly Technical Publication for
Internet and Intranet Professionals

In This Issue

From the Editor	1
A Decade of Internet Evolution.....	2
A Decade in the Life of the Internet	7
Mobile WiMAX	19
Letters to the Editor.....	36
Fragments	39

You can download IPJ
back issues and find
subscription information at:
www.cisco.com/ipj

FROM THE EDITOR

Ten years ago we published the first issue of *The Internet Protocol Journal* (IPJ). Since then, 41 issues and a total of 1,612 pages have been produced. Today, IPJ has about 37,000 subscribers all around the world. Although most of our readers prefer the paper edition, a growing number of subscribers are reading IPJ online or downloading the PDF version. This shift in reading habits may be related to the changes in technology over the last 10 years. Lower costs and higher-resolution displays and printers, as well as improvements in Internet access technologies, have made the online “experience” a lot better than in 1998.

Publishing is by no means the only area that has seen dramatic changes in the last decade. We asked Vint Cerf and Geoff Huston to reflect on Internet developments in this period, and the resulting articles, “A Decade of Internet Evolution” and “A Decade in the Life of the Internet,” are included in this issue.

Let me take this opportunity to thank all those people who have made IPJ possible. Our authors deserve a round of applause for carefully explaining both established and emerging technologies. They are assisted by an equally insightful set of reviewers and advisors who provide feedback and suggestions on every aspect of our publications process. The process itself relies heavily on two individuals: Bonnie Hupton, our copy editor, and Diane Andrada, our designer. Thanks go also to our printers and mailing and shipping providers. Last, but not least, our readers provide encouragement, suggestions, and feedback. This journal would not be what it is without them.

Because we are considering some Internet history in this issue, I would like to announce a project that takes us even further back. Before joining Cisco in 1998 I worked at the Interop Company, where I was responsible for the monthly publication of *ConneXions—The Interoperability Report*, published from 1987 through 1996. Unlike IPJ, *ConneXions* was produced in the “old-fashioned way” using various pieces of text and artwork assembled onto paste-up boards, and then photographed for subsequent plate making and offset printing. Thus no PDF files were produced at the time, but I am pleased to announce that *The Charles Babbage Institute* at the University of Minnesota has scanned the complete collection (117 issues) and it is now available at: <http://www.cbi.umn.edu/hostedpublications/Connexions/index.html>

Our final article is a look at Mobile WiMAX. WiMAX is an emerging technology that was originally designed as a fixed wireless broadband technology, a “DSL replacement,” but has evolved to support mobility.

— Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

A Decade of Internet Evolution

by Vinton G. Cerf, Google

In 1998 the Internet had about 50 million users, supported by approximately 25 million servers (Web and e-mail hosting sites, for example, but not desktops or laptops). In that same year, the *Internet Corporation for Assigned Names and Numbers* (ICANN)^[1] was created. Internet companies such as Netscape Communications, Yahoo!, eBay, and Amazon were already 3 to 4 years old and the Internet was in the middle of its so-called “dot-boom” period. Google emerged that year as a highly speculative effort to “organize the world’s information and make it accessible and useful.” Investment in anything related to the Internet was called “irrational exuberance” by the then head of the U.S. Federal Reserve Bank, Alan Greenspan.

By April 2000, the Internet boom ended—at least in the United States—and a notable decline in investment in Internet application providers and infrastructure ensued. Domino effects resulted for router vendors, Internet service providers, and application providers. An underlying demand for Internet services remained, however, and it continued to grow, in part because of the growth in the number of Internet users worldwide.

During this same period, access to the Internet began to shift from dial-up speeds (on the order of kilobits to tens of kilobits per second) to broadband speeds (often measured in megabits per second). New access technologies such as digital subscriber loops and dedicated fiber raised consumer expectations of Internet capacity, in turn triggering much interest in streaming applications such as voice and video. In some locales, consumers could obtain gigabit access to the Internet (for example, in Japan and Stockholm). In addition, mobile access increased rapidly as mobile technology spread throughout the world, especially in regions where wireline telephony had been slow to develop.

Today the Internet has an estimated 542 million servers and about 1.3 billion users. Of the estimated 3 billion mobile phones in use, about 15 percent are Internet-enabled, adding 450 million devices to the Internet. In addition, at least 1 billion personal computers are in use, a significant fraction of which also have access to the Internet. The diversity of devices and access speeds on the Internet combine to produce challenges and opportunities for Internet application providers around the world. Highly variable speeds, display areas, and physical modes of interaction create a rich but complex canvas on which to develop new Internet applications and adapt older ones.

Another well-documented but unexpected development during this same decade is the dramatic increase in user-produced content on the Internet. There is no question that users contributed strongly to the utility of the Internet as the World Wide Web made its debut in the early 1990s with a rapidly growing menu of Web pages.

But higher speeds have encouraged user-produced audio and video archives (*Napster* and *YouTube*), as well as sharing of all forms of digital content through peer-to-peer protocols. Voice over IP, once a novelty, is very common, together with video conferencing (*iChat* from Apple, for example).

Geographically indexed information has also emerged as a major resource for Internet users. In the scientific realm, *Google Earth* and *Google Maps* are frequently used to display scientific data, sensor measurements, and so on. Local consumer information is another common theme. When I found myself in the small town of Page, Arizona, looking for saffron to make paella while in a houseboat on Lake Powell, a Google search on my Blackberry quickly identified markets in the area. I called one of them and verified that it had saffron in stock. I followed the map on the Website and bought 0.06 ounces of Spanish saffron for about \$12.99. This experience reinforced my belief that having locally useful information at your fingertips no matter where you are is a powerful ally in daily living.

New business models based on the economics of digital information are also emerging. I can recall spending \$1,000 for about 10 MB of disk storage in 1979. Recently I purchased 2 TB of disk storage for about \$600. If I had tried to buy 2 TB of disk storage in 1979, it would have cost \$200 million, and probably would have outstripped the production capacity of the supplier. The cost of processing, storing, and transporting digital information has changed the cost basis for businesses that once required the physical delivery of objects containing information (books, newspapers, magazines, CDs, and DVDs). The Internet can deliver this kind of information in digital form economically—and often more quickly than physical delivery. Older businesses whose business models are based on the costs of physical delivery of information must adapt to these new economics or they may find themselves losing business to online competitors. (It is interesting to note, however, that the Netflix business, which delivers DVDs by postal mail, has a respectable data rate of about 145 kbps per DVD, assuming a 3-day delivery time and about 4.7 GB per DVD. The CEO of Netflix, Reed Hastings, told me nearly 2 years ago that he was then shipping about 1.9 million DVDs per day, for an aggregate data rate of about 275 Gbps!)

Even the media that have traditionally been delivered electronically such as telephony, television, and radio are being changed by digital technology and the Internet. These media can now be delivered from countless sources to equally countless destinations over the Internet. It is common to think of these media as being delivered in streaming modes (that is, packets delivered in real time), but this need not be the case for material that has been prerecorded. Users of iPods have already discovered that they can download music faster than they can listen to it.

With gigabit access to the Internet, one could download an hour's worth of conventional video in about 16 seconds. This fact certainly changes my understanding of "video on demand" from a streaming delivery to a file transfer. The latter is much easier on the Internet because one is not concerned about packet inter-arrival times (jitter), loss, or even orderly delivery because the packets can be reordered and retransmitted during the file transfer. I am told that about 10 hours of video are being uploaded to YouTube per second.

The battles over *Quality of Service* (QoS) are probably not over yet either. Services such as *Skype* and applications such as iChat from Apple demonstrate the feasibility of credible, real-time audio and video conferencing on the "best-efforts" public Internet. I have been surprised by the quality that is possible when both parties have reasonably high-capacity access to the Internet.

Technorati is said to be tracking on the order of 112 million blogs, and the *China Internet Network Information Center* (CNNIC) estimates 72 million Chinese blogs that are probably in addition to those tracked by Technorati. Adding to these are billions of Web pages and, perhaps even more significant, an unknown amount of information online in the form of large databases. The latter are not indexed in the same way that Web pages can be, but probably contain more information. Think about high-energy physics information, images from the Hubble and other telescopes, radio telescope data including the *Search for Extra-Terrestrial Intelligence* (SETI)^[2], and you quickly conclude that our modern society is awash in digital information.

It seems fair to ask how long accessibility of this information is likely to continue. By this question I do not mean that it may be lost from the Internet but, rather, that we may lose the ability to interpret it. I have already encountered such problems with image files whose formats are old and whose interpretation by newer software may not be possible. Similarly, I have ASCII text files from more than 20 years ago that I can still read, but I no longer have operating software that can interpret the formatting instructions to produce a nicely formatted page. I sometimes think of this problem as the "year 3000" problem: It is the year 3000 and I have just finished a Google search and found a PowerPoint 1997 file. Assuming I am running Windows 3000, it is a fair question whether the format of this file will still be interpretable. This problem would arise even if I were using open-source software. It seems unlikely that application software will last 1000 years in the normal course of events unless we deliberately take steps to preserve our ability to interpret digital content. Absent such actions, we will find ourselves awash in a sea of rotting bits whose meaning has long since been lost.

This problem is not trivial because questions will arise about intellectual property protection of the application, and even the operating system software involved. If a company goes out of business or asserts that it will no longer support a particular version of an application or operating system, do we need new regulations that require this software to be available on the public Internet in some way?

Even if we have skirted this problem in the past by rendering information into printed form, or microfilm, the complexity of digital objects is increasing. Consider spreadsheets or other complex objects that really cannot be fully “rendered” without the assistance of application software. So it will not be adequate simply to print or render information in other long-lived media formats. We really will need to preserve our ability to read and interpret bits.

The year 2008 also marks the tenth anniversary of a project that started at the U.S. Jet Propulsion Laboratory: *The Interplanetary Internet*. This effort began as a protocol design exercise to see what would have to change to make Internet-like capability available to manned and robotic spacecraft. The idea was to develop networking technology that would provide to the space exploration field the kind of rich and interoperable networking between spacecraft of any (Earth) origin that we enjoy between devices on the Internet.

The design team quickly recognized that the standard TCP/IP protocols would not overcome some of the long delays and disruptions to be expected in deep space communication. A new set of protocols evolved that could operate above the conventional Internet or on underlying transport protocols more suited to long delays and disruption. Called “delay and disruption tolerant networking”^[3, 4] or DTN, this suite of protocols is layered in the same abstract way as the Internet. The Interplanetary system could be thought of as a network of Internets, although it is not constrained to use conventional Internet protocols. The analog of IP is called the *Bundle Protocol*^[5], and this protocol can run above TCP or the *User Datagram Protocol* (UDP) or the new *Licklider Transport Protocol* (for deep space application). Ironically, the DTN protocol suite has also proven to be useful for terrestrial applications in which delay and disruption are common: tactical military communication and civilian mobile communication.

After 10 years of work, the DTN system will be tested onboard the Deep Impact mission platform late in 2008 as part of a program to qualify the new technology for use in future space missions. It is hoped that this protocol suite can be standardized for use by any of the world’s space agencies so that spacecraft from any country will be interoperable with spacecraft of other countries and available to support new missions if they are still operational and have completed their primary missions. Such a situation already exists on Mars, where the Rovers are using previously launched orbital satellites to relay information to Earth’s Deep Space Network using store-and-forward techniques like those common to the Internet.

The Internet has gone from dial-up to deep space in just the past 10 years. One can only begin to speculate about its application and condition 10 years hence. We will all have to keep our subscriptions to *The Internet Protocol Journal* to find out!

References

- [1] Cerf, V., "Looking Toward the Future," *The Internet Protocol Journal*, Volume 10, No. 4, December 2007.
- [2] <http://www.seti.org>
- [3] <http://www.dtnrg.org/wiki>
- [4] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, "Delay-Tolerant Networking Architecture," RFC 4838, April 2007.
- [5] Scott, K., and S. Burleigh, "Bundle Protocol Specification," RFC 5050, November 2007.

VINTON G. CERF is vice president and chief Internet evangelist for Google. Cerf served as a senior vice president of MCI from 1994 through 2005. Widely known as one of the "Fathers of the Internet," Cerf is the co-designer of the TCP/IP protocols and the architecture of the Internet. He received the U.S. National Medal of Technology in 1997 and the 2004 ACM Alan M. Turing award. In November 2005, he was awarded the Presidential Medal of Freedom. Cerf served as chairman of the board of the Internet Corporation for Assigned Names and Numbers (ICANN) from 2000 through 2007 and was founding president of the Internet Society. He is a Fellow of the IEEE, ACM, the American Association for the Advancement of Science, the American Academy of Arts and Sciences, the International Engineering Consortium, the Computer History Museum, and the National Academy of Engineering. He is an honorary Freeman of the City of London. Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA. E-mail: vint@google.com

A Decade in the Life of the Internet

by Geoff Huston, APNIC

The evolutionary path of any technology can often take strange and unanticipated turns and twists. At some points simplicity and minimalism can be replaced by complexity and ornamentation, while at other times a dramatic cut-through exposes the core concepts of the technology and removes layers of superfluous additions. The technical evolution of the Internet appears to be no exception, and contains these same forms of unanticipated turns and twists.

This article presents a personal perspective of the evolution of the Internet over the last decade, highlighting my impressions of what has worked, what has not, and what has changed over this period. It has been an extraordinary decade for the Internet, encompassing a boom and a bust that would rate among history's best, a comprehensive restructuring of the communications industry, and a set of changes that have altered the way in which each of us now works and plays. And the Internet has even added a few new words to the language on the way.

Rather than offer a set of random observations, I will use the Internet Protocol model as a template, starting with the underlying transmission media, then looking at the internetwork layer, the transport layer, then applications and services, and, finally looking at the business of the Internet.

The Transmission Media Layer

It seems like it was in an entirely different lifetime, but the *Internet Service Provider* (ISP) business of 1998 was still centrally involved in the technology of dial-up modems. The state-of-the-art of modem speed had been continually refined from 9,600 bps to 14.4 kbps, to 28 kbps, to finally, 56 kbps, squeezing every last bit out the phase amplitude space contained in an analogue 3-KHz voice circuit. Modems were the bane of an ISP's life. They were capricious, constantly being superseded by the next technical refinement, unreliable, difficult for customers to use, and they were just slow. Almost everything else on the Internet was tailored to download reasonably quickly over a modem connection. Webpages were carefully tailored with compressed images, and plaintext was the dominant medium as a consequence.

Not all forms of Internet access were dial-up. ISDN was used in some places, but it was never cheap enough to take over as the ubiquitous access method. There were also access services based on *Frame Relay*, X.25, and various forms of digital data services. At the high end of the speed spectrum were T1 access circuits with 1.5-Mbps clocking, and T3 circuits clocked at 45 Mbps.

ISPs leased circuits from a telephony company (telco). In 1998 the ISP industry was undergoing a transition of its trunk IP infrastructure from T1 circuits to T3 circuits. It was not going to stop here, but squeezing even more capacity from the network was proving to be a challenge. Deployment of 622-Mbps IP circuits occurred, although many of these were constructed using 155-Mbps *Asynchronous Transfer Mode* (ATM) circuits using router load balancing to share the IP load over four of these circuits in parallel. Gigabit circuits were just beginning, and the initial tests of IP over 2.5-Gbps *Synchronous Digital Hierarchy* (SDH) circuits began in 1998.

In some ways 1998 was a pivotal year for IP transmission. Until this time IP was still just another application that was positioned as just another customer of the telco's switched-circuit infrastructure that was constructed primarily to support telephony. From the analogue voice circuits to the 64K digital circuit through to the trunk bearers, IP had been running on top of the voice network. By 1998 things were changing. The Internet had started to make ever larger demands on transmission capacity, and the factor accelerating further growth in the network was now not voice, but data. It made little sense to provision an ever larger voice-based switching infrastructure just to repackage it as IP, and by 1998 the industry was starting to consider just what an all-IP high-speed network would look like, from the photon all the way through to the application.

At the same time the fiber-optic systems were changing with the introduction of *Wavelength-Division Multiplexing* (WDM). Older fiber equipment with electro-optical repeaters and *Plesiochronous Digital Hierarchy* (PDH) multiplexers allowed a single fiber pair to carry around 560 Mbps of data. WDM allowed a fiber pair to carry multiple channels of data using different wavelengths, with each channel supporting a data rate of up to 10 Gbps. Channel capacity in a fiber strand is between 40 to 160 channels using *Dense WDM* (DWDM). Combined with the use of all-optical amplifiers, the most remarkable part of this entire evolution in fiber systems is that a Tbps cable system can be constructed today for much the same cost as a 560-Mbps cable system of the mid-1990s. The factor that accelerated deployment of these high-capacity fiber systems was never based on expansion of telephony, because the explosive growth of the industry was all about IP. So it came as no surprise that at the same time as the demand for IP transmission was increasing there was a shift in the transmission model, where instead of plugging routers into telco switching gear and using virtual point-to-point circuits for IP, we started to plug routers into wavelengths of the DWDM equipment and operate all-IP networks in the core of the Internet.

The evolution of access networks has seen a shift away from modems to numerous digital access methods, including DSL, cable modems, and high-speed wireless services. The copper pair of the telco network has proved surprisingly resilient, and DSL has achieved speeds of tens of megabits per second through this network, with the prospect of hundred-megabit systems appearing soon.

So, in terms of transmission, the last 10 years has seen the network migrate from an overlay system of kilobit-per-second access with multimegabit trunks operating as a customer of the telco switched network to a comprehensive IP network with access of megabits per second with multigigabit trunks, or a thousandfold increase in basic network capacity in that period.

The demand of the Internet for capacity continues, and we are now seeing work on standardizing 40- and 100-Gbps transmission systems in the IEEE; the prospect of terabit transmissions is now taking shape for the Internet.

The Internet Layer

If transmission has seen dramatic changes in the past decade, then what has happened at the IP layer over the same period?

The glib answer is “absolutely nothing!” But that answer would be ignoring a large amount of activity in this area. We have tried to change many parts of IP in the past decade, but, interestingly, none of the proposed changes has managed to gain any significant traction in the network, and IP today is largely no different from IP of a decade ago. *Mobility*^[1], *Multicast*^[2], and *IP Security* (IPSec)^[3] remain poised in the wings, still awaiting adoption by the Internet mainstream.

Quality of Service (QoS) was a “hot” topic in 1998, and it involved the search for a reasonable way for some packets to take the fast path while others took a more leisurely way through the network. We experimented with various forms of signaling, packet classifiers, queue-management algorithms, and interpretations of the *Type of Service* bits in the IPv4 packet header, and we explored the QoS architectures of *Integrated and Differentiated Services* in great detail. However, QoS never managed to achieve wide acceptance in mainstream Internet service environments. In this case the Internet took a simpler direction: In response to not enough network capacity, the alternate approach to installing additional mechanisms in the network—in the host protocol stack and even in the application in order to ration the capacity you have—is to simply expand the network to meet the total level of demand. So far the simple approach has prevailed in the network, and QoS remains largely unused^[4].

We have experimented with putting circuits back into the IP architecture in various ways, most notably with the *Multiprotocol Label Switching* (MPLS) technology^[5]. This technology used the label-swapping approach used in X.25, Frame Relay, and ATM virtual circuit switching systems; it created a collection of virtual paths from each network ingress to each network egress. The idea was that in the interior of the network you no longer needed to load up a complete routing table into each switching element, and instead of performing destination-address lookup you could perform a much smaller, and hopefully faster, label lookup.

This process did not eventuate, and switching packets using the 32-bit destination address continued to present much the same level of cost-efficiency at the hardware level as virtual circuit label switching. When you add the additional overhead of an additional level of indirection in terms of operational management of MPLS networks, MPLS became another technology that so far has not managed to achieve traction in mainstream Internet networks. However, MPLS is by no means a dormant technology, and one place where MPLS has enjoyed considerable deployment is in the corporate service sector where many *Virtual Private Networks*^[6] are constructed using MPLS as the core technology, steadily replacing a raft of traditional private data systems that used X.25, Frame Relay, ATM, *Switched Multimegabit Data Service* (SMDS), and switched Ethernet.

Of course one change at the IP level of the protocol stack that was intended in the past decade but has not occurred is *IP Version 6*^[7]. In 1998 we were forecasting that we would have consumed all the remaining unallocated IPv4 addresses by around 2008. We were saying at the time that, because we had completed the technical specification of IPv6, the next step was that of deployment and transition. There was no particular sense of urgency, and the comfortable expectation was that with a decade to go we did not need to raise any alarms. And this plan has worked, to some extent, in that today's popular desktop operating systems of Windows, MacOS, and UNIX all have IPv6 support. But other parts of this transition have been painfully slow. It was only a few months ago that the root of the *Domain Name System* (DNS) was able to answer queries using the IPv6 protocol as transport, and provide the IPv6 addresses of the root nameservers. Very few mainstream services are configured in a dual-stack fashion, and the prevailing view is still that the case for IPv6 deployment has not yet reached the necessary threshold. Usage measurements for IPv6 point to a level of deployment of around one-thousandth of the IPv4 network, and, perhaps more worrisome, this metric has not changed to any appreciable level in the past 4 years. So what about that projection of IPv4 unallocated pool exhaustion by 2008? How urgent is IPv6 now? The good news is that the *Internet Assigned Numbers Authority* (IANA) still has some 16 percent of the address space in its unallocated pool, so IPv4 address exhaustion is unlikely to occur this year. The bad news is that the global consumption rate of IP addresses is now at a level such that the remaining address pool can fuel the Internet for less than a further 3 years, and the exhaustion prediction is now sometime around 2010 to 2011.

So why have we not deployed IPv6 more seriously yet? And if we are not going to deploy IPv6, then what is the alternative? Of all the technical refinements to IP that have occurred, one that received little fanfare when it was first published has enjoyed massive deployment over the past decade, and that is the technology of *Network Address Translation* (NAT)^[8]. Today NAT devices are ubiquitous. It seems that every home access unit, every corporate firewall, every data center, and every service includes a NAT device.

One measure of the ubiquity of NATs is the transformation that has occurred in the application space. By 2008 applications have either adopted a strict client-server approach, where the client always initiates the network transaction, or were forced down a more complex path. Where there is some form of peer interaction, applications are now equipped with additional capabilities, including NAT behavior discovery, NAT binding management, application-level name spaces, and multiparty rendezvous mechanisms, all required to allow the application to function across NATs. So far we have managed to offload the problem of looming address scarcity in the Internet onto NATs, and the really significant change that has occurred in the past decade at the IP level is the default assumption about the semantics of an IP address. An IP address is no longer synonymous with the persistent identity of the remote party that anyone can use to initiate a communication, but a temporary token to allow a single transaction to complete. As a consequence, most Internet services have retreated into data centers and the business of hosting services has thrived. And the change that would have preserved the coherent end-to-end architecture of the Internet IP layer, namely IPv6, is still waiting for wide-scale deployment.

The next few years promise to be “interesting” in every form of meaning of the word. The exhaustion of the remaining IPv4 address pool is imminent, and if we are going to substitute IPv6 in place of IPv4, then we simply do not have enough time to achieve this substitution before the remaining IPv4 address pool is depleted. And although so far NATs have conveniently pushed the problem of increasing address scarcity off the network and over to the edge devices and onto applications, it is not clear that this approach can sustain an ever-growing Internet indefinitely. We have yet to understand just what a “carrier-grade NAT” might be, or whether it can even work in any useful manner. NATs were an accidental addition to the Internet, and their role in the coming years is unclear.

The early 1990s saw a flurry of activity in the routing space, and protocols were quickly developed and deployed. By 1998 the “standard” Internet environment involved the use of either *Intermediate System-to-Intermediate System* (IS-IS) or *Open Shortest Path First* (OSPF) as large-scale interior routing protocols and *Border Gateway Protocol 4* (BGP4) as the interdomain routing protocol^[9]. This picture has remained constant over the past decade. In some ways it is reassuring to see a technology that is capable of sustaining a quite dramatic growth rate, but perhaps that is not quite the complete picture.

We never quite completed the specification for the next interdomain routing protocol, and BGP4 is now showing signs of stress^[10]. The pool of *Autonomous System* (AS) numbers is forecast to run out early in 2011, and by then we need to have fielded a new variant of BGP that can operate with a much larger pool of AS numbers^[11].

Fortunately the technology development has been completed and an approach that allows incremental deployment has been devised, so this transition is not quite the traumatic transition that is associated with IPv6. But deployment is slow, and of the current level of adoption of the larger AS number set is, oddly enough, comparable to IPv6, at a level of around one-thousandth of the total AS number pool. The routing system has also been growing inexorably, and the capability of switching systems to cope with ever larger routing tables while at the same time offering continual improvements in cost-efficiencies is now looking less certain. So, once again we appear to be examining routing protocol theory and practice, and looking at alternate approaches to routing that can offer superior scaling properties to BGP for the future.

No listing of the major highlights in IP over the past decade would be complete without some mention of the perennial issue of *location* and *identity*.^[25] One of the original simplifications in the IP architecture was to place the semantics of identity, location, and forwarding into an IP address. Although that process has proved phenomenally effective in terms of simplicity of applications and simplicity of IP networks, it has posed some serious challenges with regard to mobility, routing, and network management. Each of these aspects of the Internet would benefit considerably if the Internet architecture allowed identity to be distinct from location. Numerous efforts have been directed at this problem over the past decade, particularly in IPv6, but so far we really have not arrived at an approach that feels truly comfortable in the context of IP.

So although it is possible to observe that not much has happened at the IP level in the past decade that is deployed in the Internet—and IP is still IP—there is still a considerable agenda to tackle at the Internet layer.

The Transport Layer

A decade ago, in 1998, the transport layer of the IP architecture consisted of the *User Datagram Protocol* (UDP) and TCP, and the network usage pattern was around 95-percent TCP and 5-percent UDP. Here, as well, not much has changed in the intervening 10 years.

We have developed two new transport protocols, the *Datagram Congestion Control Protocol* (DCCP) and the *Stream Control Transmission Protocol* (SCTP)^[12], which can be regarded as refinements of TCP to cover flow control for datagram streams in the case of DCCP and flow control over multiple reliable streams in the case of SCTP. However, in a world of transport-aware middleware that is the Internet today, the level of capability to actually deploy these new protocols in the public Internet is marginal at best.

TCP has proved to be remarkably resilient over the years, but as the capacity of the network increases the ability of TCP to continue to deliver ever faster data rates over distances that span the globe is becoming a significant concern. Recent times have seen much work to devise revised TCP flow-control algorithms that still share the network fairly with other concurrent TCP sessions, yet can ramp up to multigigabit-per-second data-transfer rates and sustain those rates over extended periods^[13]. At this stage much of this work is still in the area of research and experimentation, and TCP today as deployed on the Internet is much the same as TCP of a decade ago, with perhaps a couple of notable exceptions. The latest TCP stack from Microsoft in Vista uses dynamic tuning of the Receive window, and a larger inflation factor of the Send window in congestion avoidance where there is a large bandwidth delay product, and improved loss-recovery algorithms that are particularly useful in wireless environments. Linux now includes an implementation of *Binary Increase Congestion control* (BIC), which undertakes a binomial search to reestablish a sustainable send rate. Both of these approaches can improve the performance of TCP, particularly when sending the TCP session over long distances and trying to maintain high transfer speeds.

The Application and Service Layer

This area, unlike the transport layer, has seen quite profound changes over the past decade. A decade ago the Internet was on the cusp of portal mania, where *LookSmart* was the darling of the Internet boom and everyone were all trying to promote their own favorite “one stop shop” for all their Internet needs. We were still using various forms of hand-compiled directories, and navigation of the Internet was still the subject of various courses and books.

By 1998 *AltaVista* has made its debut, and change was already evident. This change, from directories and lists to active search, completely changed the Internet. These days we simply assume that we can type any query we have into a search engine and the search machinery will deliver a set of pointers to relevant documents. Each time this process occurs our expectations about the quality and utility of search engines are reinforced, and we have moved beyond swapping URLs as pointers and simply exchange search terms as an implicit reference to the material. Content is also changing as a result, because users no longer remain on a “site” and navigate around the site. Instead users are directing the search engines, and pulling the relevant page from the target site without reference to any other material.

Another area of profound change has been the rise of active collaboration over content, best typified in wikis. *Wikipedia* is perhaps the most cited example of user-created content, but almost every other aspect of content generation is also being introduced into the active user model, including *YouTube*, *Flickr*, *Joost*, and similar content.

Underlying these changes is another significant development, namely the changes in the content economy. In 1998 content providers and ISPs were competing for user revenue. Content providers were unable to make pay per view and other forms of direct financial relationship with users work in their favor, and were arguing that ISPs should fund content, because, after all, the only reason that users paid for Internet access was because of their perceived value of the content. ISPs, on the other hand, promoted the idea that content providers were enjoying a free ride across the ISP-funded infrastructure, and content providers should contribute to network costs. The model that has gained ascendancy as a result of this unresolved tension was that of advertised-funded content services, and this model has sustained a vastly richer, larger, and more compelling content environment.

At the same time the peer-to-peer network has emerged, and from its beginnings as a music-sharing subsystem, the distributed data model of content sharing now dominates the Internet with audio, video, and large data sets now using this form of content distribution and its associated highly effective transport architecture. Various measurements of Internet traffic have placed peer-to-peer content movement at between 40 and 80 percent of the overall traffic profile of the network.

In many ways applications and services have been the high frontier of innovation in the Internet in the past decade. An entire revolution in open interconnection of content elements is embraced under the generic term *Web 2.0*, and “content” is now a very malleable concept. It is no longer the case of “my computer, my applications, and my workspace” but an emerging model where not only the workspace for each user is held in the network, but where the applications themselves are part of the network, and all are accessed through a generic browser interface.

Any summary of the evolution of the application space over the last decade would not be complete without noting that whereas in 1998 the Internet was still an application that sat on top of the network infrastructure used to support the telephone network, by 2008 voice telephony was just another application layered on the infrastructure of the Internet, and the Internet had even managed to swallow the entire telephone number space into its DNS, using an approach called *ENUM*^[14].

The Business Layer

As much as the application environment of the Internet has been wildly erratic over the past decade, the business environment has been unpredictable as well, and the list of business winners and losers includes some of the historical giants of the telephone world as well as the Internet-bred new wave of entrants.

In 1998, despite the growing momentum of public awareness, the Internet was still largely a curiosity. It was an environment inhabited by geeks, game players, and academics, whose rites of initiation were quite arcane. As a part of the data networking sector, the Internet was just one further activity among many, and the level of attention from the mainstream telco sector was still relatively small. Most Internet users were customers of independent ISPs, and the business relationship between the ISP sector and the telco was tense and acrimonious. The ISPs were seen as opportunistic leeches on the telco industry; they ordered large banks of phone lines, but never made any calls; their customers did not hang up after 3 minutes, but kept their calls open for hours or even days at a time, and they kept ordering ever larger inventories of transmission capacity, yet had business plans that made the back of an envelope look professional by comparison. The telco was unwilling to make large long-term capital investments in additional infrastructure to pander to the extravagant demands of a wildcat set of Internet speculators and their fellow travelers. The telco, on the other hand was slow, expensive, inconsistent, ill-informed, and hostile to the ISP business. The telco wanted financial settlements and bit-level accounting, whereas the ISP industry appeared to manage quite well with a far simpler system of peering and tiering that avoided putting a value on individual packets or flows^[15]. This relationship was never going to last, and it resolved itself in ways that in retrospect were quite predictable. From the telco perspective it quickly became apparent that the only reason the telco was being pushed to install additional network capacity at ever increasing rates was the requirements of the ISP sector. From the ISP perspective the only way to grow at a rate that matched customer demand was to become one's own carrier and to take over infrastructure investment. And, in various ways, both outcomes occurred. Telcos bought ISPs, and ISPs became infrastructure carriers.

All this activity generated considerable investor interest, and the rapid value escalation of the ISP industry and then the entire Internet sector generated the levels of wild-eyed optimism that are associated only with an exceptional boom. By 2000 almost anything associated with the Internet, whether it was a simple portal, a new browser development, a search engine, or an ISP, attracted investor attention, and the valuations of Internet start-ups achieved dizzying heights. Of course one of the basic lessons of economic history is that every boom has an ensuing bust, and in 2001 the Internet bust happened. The bust was as inevitable and as brutal as the preceding boom was euphoric. But, like the railway boom and bust of the 1840s, when the wreckage was cleared away, what remained was a viable—and indeed a valuable—industry.

By 2003 the era of the independent retail ISP was effectively over. ISPs still exist, but those that are not competitive carriers tend to operate as IT business consultants that provide services to niche markets. Their earlier foray in to the mass market paved the way for the economies of scale that only the carrier industry could implement on the market.

But the grander aspirations of these larger players have not been met, and effective monopoly positions in many Internet access markets have not translated to effective control over the user's experience of the Internet, or anything even close to such control. The industry was already "unbundled," with intense competition occurring at every level of the market, including content, search, applications, and hosting. The efforts of the telco sector to translate their investment into mass-market Internet access into a more comprehensive control over content and its delivery in the Internet has been continually frustrated. The content world of the Internet has been reinvigorated by the successful introduction of advertiser-funded models of content generation and delivery, and this process has been coupled with the more recent innovations of turning back to the users themselves as the source of content, so that the content world is once again the focus of a second wave of optimism, bordering on euphoria.

And Now?

It has been a revolutionary decade for us all, and in the last 10 years the Internet has directly touched the lives of almost every person on this planet. Current estimates put the number of regular Internet users at 19 percent of the world's population.

Over this decade some of our expectations were achieved and then surpassed with apparent ease, whereas others remained elusive. And some things occurred that were entirely unanticipated. At the same time very little of the Internet we have today was confidently predicted in 1998, whereas many of the problems we saw in 1998 remain problems today.

What we have today is not the technical Internet we thought we were building a decade ago. It is not a coherent end-to-end network with clear signaling across commodity packet switching fabric, but a network that is replete with all forms of active middleware^[16], from NATs to firewalls^[17] and filters, including packet shapers, torrent detectors, *Voice over IP* (VoIP) blockers, and load balancers. It is neither a secure nor a safe network, but one that includes a continual barrage on end hosts in the form of more than a million different forms of viruses^[18], worms, and assorted malware^[19], as well as a barrage on users in the form of torrents of spam^[20]. The network is a host to a litany of hostile attacks, including gigabit traffic swamping attacks, redirection, inspection, passing off, and denial-of-service attacks^[21]. The attacks are directed at links, routers^[22], the routing protocols^[23, 24], hosts, and applications. Our ability to effectively defend the network and its connected hosts continues to be, on the whole, ineffectual. Our level of interest in paying a premium to support highly secure systems still remains slight. But somehow we are not deterred by this situation. Somehow each of us has found a way to make our Internet work for us.

I am not sure that the next decade will bring the same level of intensity of structural change to the global communications sector, and perhaps that is a good thing given the collection of other challenges that are confronting us all in the coming decades. At the same time I think it would be good to believe that the past decade of development of the Internet has completely rewritten what it means to communicate, rewritten the way in which we can share our experience and knowledge, and, hopefully, rewritten the ways in which we can work together on these challenges.

References

The Internet Protocol Journal (IPJ) has published articles on all the major aspects of the technical evolution of the Internet over the past decade. To illustrate the extraordinary breadth of these articles, I have included as references here only articles that have been published in the IPJ.

- [1] Stallings, W., “Mobile IP,” *IPJ*, Volume 4, No. 2, June 2001.
- [2] Handley, M., and Crowcroft, J., “Internet Multicast Today,” *IPJ*, Volume 2, No. 4, December 1999.
- [3] Stallings, W., “IP Security,” *IPJ*, Volume 3, No. 1, March 2000.
- [4] Huston, G., “QoS — Fact or Fiction?” *IPJ*, Volume 3, No. 1, March 2000.
- [5] Stallings, W., “MPLS,” *IPJ*, Volume 4, No. 3, September 2001
- [6] Ferguson, P., and Huston, G., “What is a VPN?” *IPJ*, Volume 1, No. 1 & No. 2, June & September 1998.
- [7] Fink, R., “IPv6,” *IPJ*, Volume 2, No. 1, March 1999.
- [8] Huston, G., “Anatomy: Inside Network Address Translators,” *IPJ*, Volume 7, No. 3, September 2004.
- [9] Huston, G., “The BGP Routing Table,” *IPJ*, Volume 4, No. 1, March 2001.
- [10] Huston, G., “Scaling inter-Domain Routing,” *IPJ*, Volume 4, No. 4, December 2001.
- [11] Huston, G., “Exploring Autonomous System Numbers,” *IPJ*, Volume 9, No. 1, March 2006.
- [12] Huston, G., “The Future for TCP,” *IPJ*, Volume 3, No. 3, September 2000.
- [13] Huston, G., “Gigabit TCP,” *IPJ*, Volume 9, No. 2, June 2006.

- [14] Huston, G., "ENUM," *IPJ*, Volume 5, No. 2, June 2002.
- [15] Huston, G., "Peering and Settlements," *IPJ*, Volume 2, No. 1 & No. 2, March & June 1999.
- [16] Huston, G., "The Middleware Muddle," *IPJ*, Volume 4, No. 2, June 2001.
- [17] Avolio, F., "Firewalls and Internet Security," *IPJ*, Volume 2, No. 2, June 1999.
- [18] Fraser, B., Rogers, L., and Pesante, L., "Was the Melissa Virus So Different?" *IPJ*, Volume 2, No. 2, June 1999.
- [19] Chen, T., "Virus Trends," *IPJ*, Volume 6, No. 3, September 2003.
- [20] Crocker, D., "Challenges in Anti-Spam Efforts," *IPJ*, Volume 8, No. 4, December 2005.
- [21] Patrikakis, C., Masikos, M., and Zouraraki, O., "Distributed Denial of Service Attacks," *IPJ*, Volume 7, No. 4, December 2004.
- [22] Lonvick, C., "Securing the Infrastructure," *IPJ*, Volume 3, No. 3, September 2000.
- [23] Kent, S., "Securing BGP: S-BGP," *IPJ*, Volume 6, No. 3, September 2003.
- [24] White, R., "Securing BGP: soBGP," *IPJ*, Volume 6, No. 3, September 2003.
- [25] Meyer, D., "The Locator Identifier Separation Protocol (LISP)," *IPJ*, Volume 11, No. 1, March 2008.

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. The author of numerous Internet-related books, he is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001. E-mail: gih@apnic.net

Mobile WiMAX

by Jarno Pinola and Kostas Pentikousis, VTT Technical Research Centre of Finland

One of the technologies that can lay the foundation for the next generation (fourth generation [4G]) of mobile broadband networks is popularly known as “WiMAX.” WiMAX, *Worldwide Interoperability for Microwave Access*, is designed to deliver wireless broadband bitrates, with *Quality of Service* (QoS) guarantees for different traffic classes, robust security, and mobility. This article provides an overview of mobile WiMAX, which is based on the wireless local and *Metropolitan-Area Network* (MAN) standards IEEE 802.16-2004^[1] and 802.16e-2005^[2]. We introduce WiMAX and focus on its mobile system profile and briefly review the role of the WiMAX Forum. We summarize the critical points of the WiMAX network reference model and present the salient characteristics of the PHY and MAC layers as specified in [1] and [2]. Then we address how mobile nodes enter a WiMAX network and explain the fundamentals of mobility support in WiMAX. Finally, we briefly compare WiMAX with *High-Speed Packet Access* (HSPA), another contender for 4G.

The Role of the WiMAX Forum

The WiMAX Forum is a nonprofit organization formed in 2001 to enhance the compatibility and interoperability of equipment based on the IEEE 802.16 family of standards. The IEEE 802.16 standards provide a large set of fundamentally different options for designing a wireless broadband system, including, for example, multiple options for *Physical* (PHY) layer implementation, *Media Access Control* (MAC) architecture, frequency bands, and duplexing. So many options lead to several possible system variants, which are all compatible with the IEEE standards. Although such multiplicity allows for deployment in very diverse environments, it may spell either solely vertical, single-vendor deployments or no deployment at all, because operators do not want to be locked in with any particular implementation. Thus, a major motivation for establishing the WiMAX Forum was to develop predefined system profiles for equipment manufacturers, which include a subset of the features included in the IEEE 802.16 standards. WiMAX Forum-certified products are guaranteed to be interoperable and to support wireless broadband services from fixed to fully mobile scenarios. The aim is to enable rapid market introduction of new standard-compliant WiMAX equipment and to promote the use of the technology in different sectors.

From IEEE 802.16 to Mobile WiMAX

The IEEE 802.16 standard was originally meant to specify a fixed wireless broadband access technique for point-to-point and point-to-multipoint links. During its development, however, it was decided that mobility support should also be considered.

The WiMAX Forum defines two system profiles based on [1] and [2], called *fixed* and *mobile* system profiles, respectively. Both include mandatory and optional PHY and MAC layer features that are required from all corresponding WiMAX-certified products. Because [1] and [2] specify only the PHY and MAC layers, an end-to-end architecture specification was deemed necessary in order to enable fast growth in manufactured quantities, market share, and interoperability. In response, the WiMAX Forum established the *Network Working Group* (NWG) with the aim of developing an end-to-end network reference model architecture based on IP supporting both fixed and mobile WiMAX (refer to [3] and [4]).

In short, according to the NWG reference model, a WiMAX network is partitioned into three independent architectural components: the user equipment (also referred to as *Customer Premises Equipment* [CPE]), the *Radio Access Network* (RAN, based on IEEE 802.16), and the network providing IP connectivity with the rest of the Internet. Clearly, this model allows a single operator to freely mix and match offerings from different manufacturers for these three parts, at least after interoperable equipment becomes readily available. Furthermore, in principle, each of these components of an operational network can be deployed and managed by different service providers. This scenario makes the network architecture flexible, eases network operation and maintenance, can increase competition under certain conditions, and is conducive to new business models. For example, municipalities can venture jointly with local or national network operators to deploy WiMAX in suburban and rural areas.

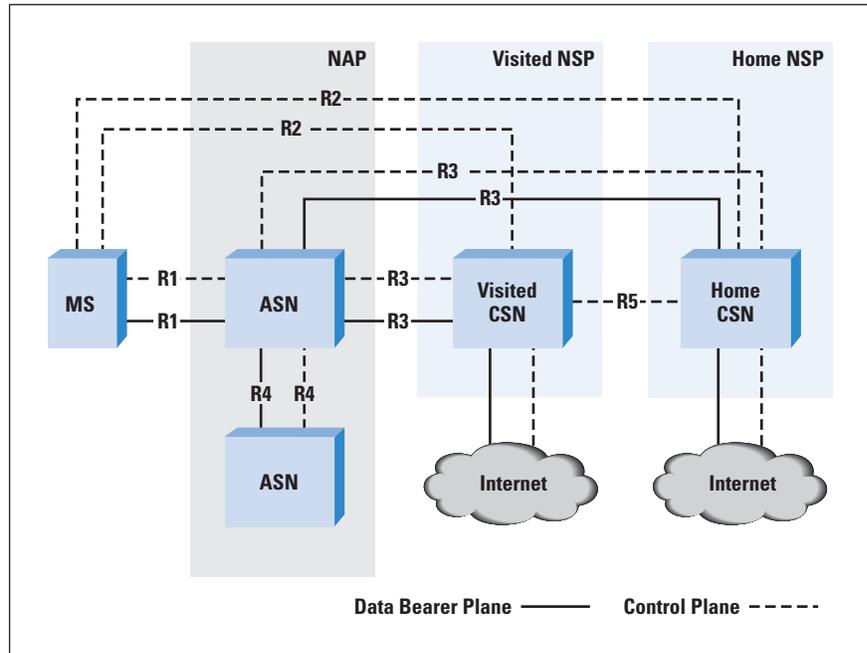
In contrast with earlier wireless data networks^[5], IP is fundamental in a WiMAX network. Indeed, IP currently plays a dominant role in the present state of the telecommunications industry. The premise is that by embracing IP, service providers and equipment manufacturers will face fewer problems when introducing WiMAX into their networks and product portfolios. Moreover, protocols standardized by the *Internet Engineering Task Force* (IETF) are preferred over proprietary solutions and are adopted as extensively as possible in the reference model.

Mobile WiMAX Network Reference Model

The WiMAX Forum NWG network reference model defines three basic architectural entities: the *Mobile Station* (MS), the *Access Service Network* (ASN), and the *Connectivity Service Network* (CSN). The role of the MS is to provide user access to the WiMAX network. The ASN is the Radio Access Network and is formed by numerous *Base Stations* (BSs) and *ASN Gateways* (ASN-GWs), managed by a *Network Access Provider* (NAP). CSN is the network entity providing IP connectivity to the WiMAX radio equipment, including all the IP core network functions required for internetworking with the rest of the world. CSNs are maintained by *Network Service Providers* (NSPs).

The ASN and CSN are further broken up into smaller functional entities, which communicate with each other using standardized interfaces called *reference points*. These reference points guarantee that a certain set of protocols and procedures are always supported and can function irrespective of the underlying hardware. The currently defined reference points are used for different control and management purposes, as well as for data bearing between the network entities. Figure 1 illustrates the network reference model and the main reference points.

Figure 1: WiMAX Forum NWG Network Reference Model



The reference points are defined as follows in [3]: Reference point R1 consists of protocols and procedures compliant to [1], [2], and [6]. R1 implements the specifications of the air interface between the MS and the BS. R2, an interface between the MS and a CSN, is used solely for management purposes, including mobility management. R3 serves the same purpose between an ASN and a CSN, and R4 is used for micromobility management between two ASNs. R5 enables interworking between two CSNs for macromobility management.

In addition to reference points R1–R5, another three intra-ASN reference points are defined (not illustrated in Figure 1). R6, which consists of a set of control- and bearer-plane protocols for BS and ASN-GW communication, controls the data path and MS mobility events between these two ASN entities. R7 is an optional set of protocols used for coordinating R6 functions. Finally, R8 consists of bearer-plane protocols that enable data transfer between the base stations involved in a handover (also called *handoff*).

With respect to mobility, the reference model considers two different scenarios called *ASN-anchored mobility* and *CSN-anchored mobility*. ASN-anchored mobility (or intra-ASN mobility, or micromobility) management is employed when MS handovers occur from one BS to another, and both are controlled by the same ASN-GW. On the other hand, CSN-anchored mobility (or inter-ASN mobility, or macromobility) management is employed when MS movement dictates a handover from the currently serving BS to another one that is in a different subnetwork, controlled by a different ASN-GW. In the ASN-anchored case, handovers are managed solely by the MS and the ASN. In the CSN-anchored case, both ASN and CSN entities are engaged in mobility management.

Typically, ASN-anchored mobility procedures take precedence and CSN-anchored mobility management is employed only if necessary. Because ASN-anchored mobility takes place inside a single ASN, it does not change the MS network layer (IP) configuration. Three different functions are specified for ASN-anchored mobility management, all considered peer-to-peer interactions between different architectural entities:

- The *handoff* (HO) function controls the handover decision operation and handover signaling. The HO function supports mobile- and network-initiated handovers and, additionally, it may support *Fast Base Station Switching* (FBSS) or *Macro Diversity Handover* (MDHO)^[2].
- The *Data Path* (DP) function manages the data path setup and data packet transmission between two functional entities.
- The context function addresses the exchanges required in order to retrieve or set up any state in the network elements.

On the other hand, when MS movement necessitates CSN-anchored mobility management, the MS IP layer configuration changes as a result of the handover. In this case, mobility management is based on *Mobile IPv4* (MIPv4)^[7] or *Mobile IPv6* (MIPv6)^[8], if the MS supports it. Alternatively, the reference model adopts *Proxy MIP* (PMIP)^[9] to handle the handover. In PMIP, the MIP function is moved from the MS to a network instance called a *PMIPv4 client*, which takes care of all MIP signaling on behalf of the MS. Support for PMIP is specified only for MIPv4 in [3] and [4]. Note that in a handover from one ASN to another, MIP is used to complement ASN-anchored mobility management. The latter is still necessary to control the link-layer handover procedures. That is, after the micromobility handover is successfully completed, MIP independently takes care of the macromobility handover, that is, establishes communication paths between the new ASN-GW and the CSN. CSN-anchored mobility handovers are always network-initiated.

By embracing IETF protocols and providing an end-to-end architecture with independent functional entities, the WiMAX Forum NWG network reference model provides a clear framework for the application developers to work in. The model provides only operational requirements and does not prescribe particular technical solutions to realize them, allowing for proprietary yet standards-compliant implementations and enabling technical competition between different manufacturers.

Before examining mobility support in WiMAX, we review the basics of the IEEE 802.16 PHY and MAC layers.

OFDM and OFDMA

IEEE 802.16 and thus WiMAX adopted *Orthogonal Frequency Division Multiplexing* (OFDM), a multicarrier modulation scheme, as its PHY layer. In OFDM, the available bandwidth is divided into several parallel orthogonal subcarriers with lower bandwidth. A wideband channel is defined as a group of adjacent narrowband channels: a high-bitrate data stream is divided into these subcarriers and multiple narrowband data streams are transmitted over the air. Because the data symbol duration is inversely proportional to bitrate, the transmitted symbol duration is increased and the level of *Inter-Symbol Interference* (ISI) can be reduced. ISI is caused by multipath propagation in the wireless communication medium, where the transmitted data symbols can arrive at the receiver through different propagation routes because of reflections from buildings in urban areas and from hills and trees in rural areas. OFDM also uses guard intervals between successive data symbols and cyclic prefixes in order to decrease the effect of ISI even more.

One reason for the wide adoption of OFDM in modern broadband communication systems is its hardware implementation simplicity. OFDM signals can be formed and processed using *Inverse Fast Fourier Transform* (IFFT) and *Fast Fourier Transform* (FFT), at the transmitter and receiver, respectively, and both transforms can be implemented directly in hardware for higher performance. OFDM bodes well for mobile broadband systems through frequency diversity and adaptivity in both modulation and channel coding. By using *Adaptive Modulation and Coding* (AMC), the end-to-end quality deterioration due to the excess delays and deep fading conditions caused by mobility can be prevented, or at least diminished.

OFDM can also be used as a multiaccess scheme by having subcarriers grouped into subchannels, which can be assigned to different users contending for the data link. Each subchannel can contain a different number of subcarriers, and by altering the subcarrier group sizes and observing the channel conditions, it is possible to use differentiation in the channel allocation for different users.

This technique of using OFDM as a multiaccess scheme is called *Orthogonal Frequency Division Multiple Access* (OFDMA). Mobile WiMAX uses OFDMA as its PHY layer instead of plain OFDM, and subchannelization to both uplink and downlink transmissions is possible.

In OFDMA, the subcarriers assigned to subchannels can be either concurrent or taken from different regions of the total bandwidth. Both of these allocation schemes have advantages. When subcarriers assigned to one subchannel are distributed over the available bandwidth, frequency diversity can be attained. In mobile systems this diversity is advantageous because it can be used to make the transmission link more resistant against fast fading. A subchannelization scheme based on dispersed subcarrier allocation to subchannels, called *Partial Usage of Subcarriers* (PUSC), is mandatory in all mobile WiMAX implementations.

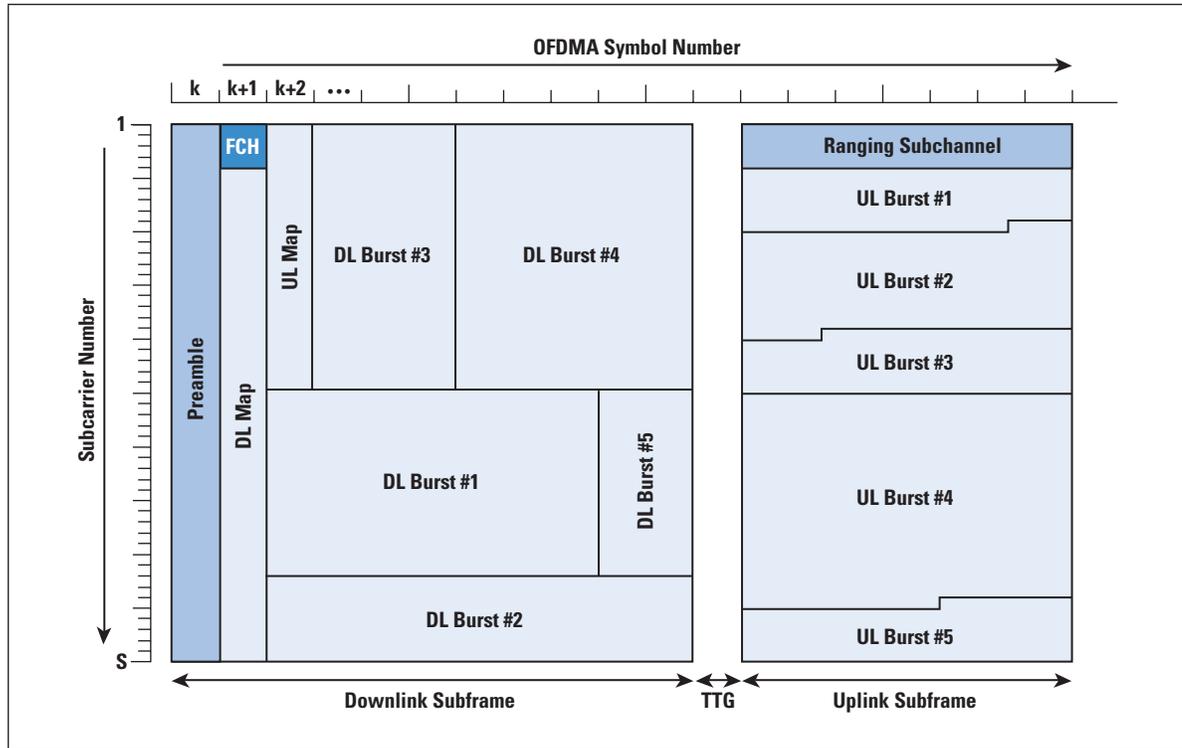
WiMAX systems can use *Time-Division Duplexing* (TDD) or *Frequency-Division Duplexing* (FDD) when allocating air interface resources to users. In TDD, the uplink and downlink transmissions are done over the same carrier frequencies and the separation between the transmission directions is done by assigning time slots, in which the transmission to one direction at a time is scheduled. In FDD, uplink and downlink transmissions are done simultaneously over different carrier frequencies.

Commonly used in mobile WiMAX equipment, TDD allows more flexible sharing of the available bandwidth between the uplink and downlink transmissions. On balance, TDD requires synchronization between multiple adjacent base stations so that transmissions in neighboring cells do not interfere with each other. A TDD frame (Figure 2) is divided into two subframes: first comes a downlink frame and after a short guard interval, called the *Transmit/Receive Transition Gap* (TTG), an uplink frame follows in the same frequency band. Each downlink subframe starts with a preamble, which is used for synchronization and channel estimation. To enhance tolerance against mobility-inflicted channel impairments, WiMAX allows optional support for a more frequent preamble repetition during transmission. In the uplink, short preambles, also called *midambles*, can be used after 8, 16, or 32 OFDM symbols, and in the downlink, short preambles in front of every data burst can be used. After the preamble comes a *Frame Control Header* (FCH), which consists of uplink and downlink *Media Access Protocol* (MAP) messages, which inform users about their transmission parameters.

Flexible data multiplexing from different users into one OFDM or OFDMA frame is also supported, as illustrated in Figure 2. Both uplink and downlink subframes can include data bursts of different types from multiple users, and they can be of variable length.

A small portion of the uplink subframe is reserved for transmission parameter adjustment and bandwidth request purposes. Moreover, small amounts of user data can be sent in this portion of the uplink subframe. The total OFDM frame size can range between 2.5 and 20 ms, but the initially supported frame size in present WiMAX equipment is 5 ms.

Figure 2: An example of a WiMAX OFDMA Frame

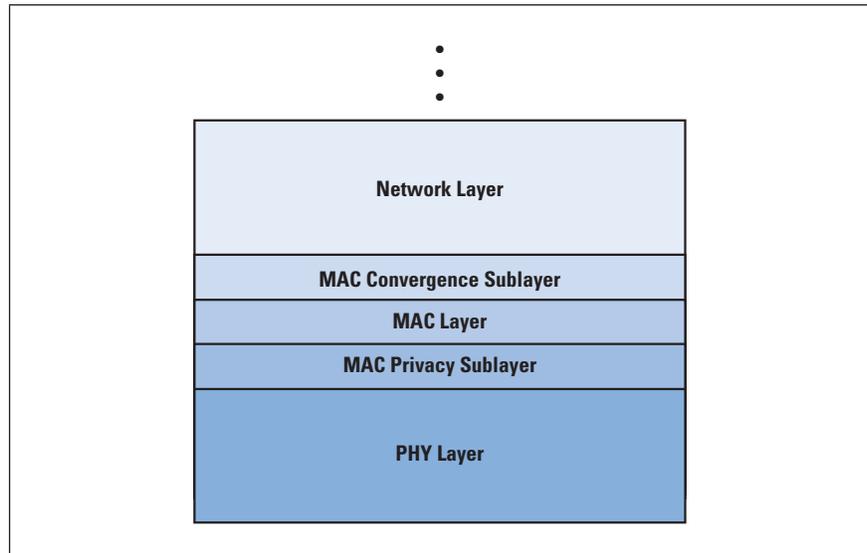


Media Access Control

The MAC layer is primarily an adaptation layer between the PHY layer and the upper layers. Its most important task, when transmitting data, is to receive *MAC Service Data Units* (MSDUs) from the layer above, aggregate and encapsulate them into *MAC Protocol Data Units* (MPDUs), and pass them down to the OFDM or OFDMA PHY layer for transmission. When data is received, the MAC layer takes MPDUs from the PHY layer, decapsulates and reorganizes them into MSDUs, and passes them on to the upper-layer protocols.

An additional layer between the MAC and upper protocol layers called the *Convergence Sublayer* (CS) is also defined in [1] and [2] and illustrated in Figure 3. For the upper layers, CS functions as an interface to the MAC layer. Even though in principle a CS is presented for a variety of different protocols, currently [3] and [4] support CS only for IP and Ethernet. Other protocols can, of course, use these CSs through encapsulation. The CS may also support upper-protocol header compression.

Figure 3: WiMAX Protocol Stack



Similarly with the PHY layer, shown in Figure 3, the MAC layer allows flexible allocation of transmission capacity to different users. Variably sized MPDUs from different flows can be included into one data burst before being handed over to the PHY layer for transmission. Multiple small MSDUs can be aggregated into one MPDU and, conversely, one big MSDU can be fragmented into multiple small ones in order to further enhance system performance. For example, by bundling up several MPDUs or MSDUs, the PHY and MAC layer header overheads, respectively, can be reduced.

It is important to remember that the BS MAC layer manages bandwidth allocation for both uplink and downlink transmissions. The BS assigns bandwidth for the downlink transmission according to incoming network traffic. For the uplink transmission, bandwidth is allocated based on the requests received from the MS. Because basically all connections are controlled by the BS, QoS can be efficiently implemented into WiMAX equipment. Currently, the MAC layer of a mobile WiMAX BS should include support for five different QoS classes, briefly summarized in Table 1.

Table 1: Mobile WiMAX QoS Classes

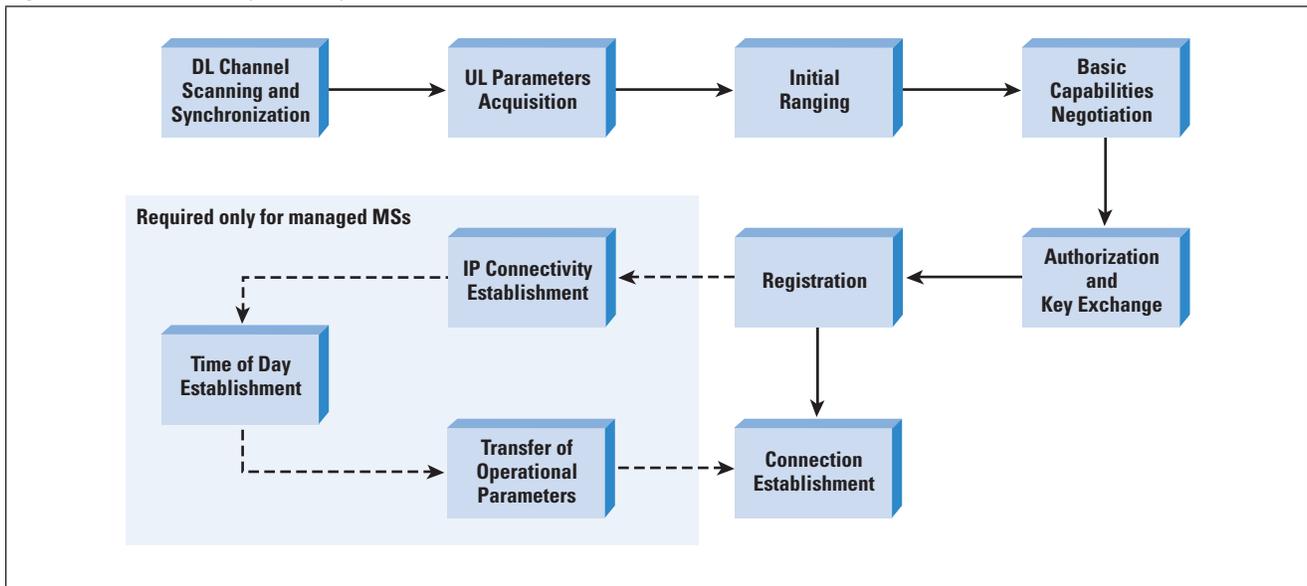
QoS Class	Supported Service	Example Application
Unsolicited Grant Services (UGS)	Latency- and jitter-sensitive applications with fixed-size data packets at Constant Bitrate (CBR)	Voice over IP (VoIP) without silence suppression
Real-Time Variable Rate (RT-VR)	Real-time applications with variable-size data packet bursts	Video and audio streaming
Non-Real-Time Polling Services (nrtPS)	Delay-tolerant applications with variable-size data packets and guaranteed bitrate demands	File transfers
Extended Real-Time Variable Rate (ERT-VR)	Real-time applications with Variable Bitrate (VBR) data streams and guaranteed bitrate and delay demands	VoIP with silence suppression
Best Effort (BE)	Data streams with no minimum service-level demands	Web browsing, instant messaging, and data transfer

Prior to any data transmission over a WiMAX link, the MS and the BS must form a unidirectional connection between their respective MAC layers. A unique identifier, called *Connection Identifier (CID)*, is assigned to each uplink and downlink connection pair. The CID serves as a temporary address for the transmitted data packets over the WiMAX link. Another identifier, called *Service Flow Identifier (SFID)*, is assigned by the BS to unidirectional packet flows with the same QoS parameters, that is, service flows. The BS also handles the mapping of SFIDs to CIDs in the QoS control process. Note that the MAC layer incorporates sophisticated power-management techniques and robust, state-of-the-art security features, but these features are out of scope for this article.

Network Entry and Reentry

Figure 4 illustrates the basic steps that every MS must go through when entering or reentering a WiMAX network. First, a MS scans the downlink channel and synchronizes with the BS, after which the MS acquires the transmit parameters for the uplink transmission from the BS *Uplink Channel Descriptor (UCD)* message and performs initial ranging, hence acquiring the correct timing offset and power adjustments. A MS extracts an initial ranging-interval time slot from an uplink MAP message. If a MS cannot complete the initial ranging successfully, it must start scanning for a new downlink channel.

Figure 4: Network Entry/Reentry Procedure



The basic capabilities negotiation process starts when the MS sends a message containing its capabilities to the BS; the BS responds with a message containing the capabilities it has in common with the MS. If *Privacy Key Management (PKM)* is enabled at both the MS and the BS, the next step is to perform the authorization and key-exchange procedure, so that the MS can register with the network. The BS sends back a registration response message that contains the secondary management CID, if the MS is managed.

After a managed MS obtains this secondary management CID, it becomes “manageable.” The successful reception of the registration response message is a prerequisite for any MS in order to be able to transmit to and receive from the network.

When a managed MS enters the network, the next step is to establish IP connectivity by using the assigned secondary management connection and by either invoking the *Dynamic Host Configuration Protocol* (DHCP)^[10] or DHCPv6^[11], or using the IPv6 stateless address autoconfiguration^[12], depending on the information provided by the BS registration response message. If the MS uses MIPv4 or MIPv6, it can secure its address by using the secondary management connection with MIP. The establishment of IP connectivity and time of day, as well as the transfer of the operational parameters, are needed only for managed MSs. These parameters can be managed with IP management messages through a secondary management connection, for example, by using the DHCP, *Trivial File Transfer Protocol* (TFTP)^[13], or *Simple Network Management Protocol* (SNMP)^[14]. These additional steps during network entry are necessary for the operation of the IP management protocols.

If DHCP is used to establish IP connectivity, a managed MS must also establish the time of day so that the management system can time-stamp certain events. Both the MS and the BS must be set at the same time of day, with an accuracy of the nearest second. The time of day is retrieved using the secondary management connection with the *Time Protocol*^[15]. The current time is formed by combining the time retrieved from the server with the time offset extracted from the DHCP reply message. Although the time of day is not needed for the registration to complete successfully, it is required in order to keep the connection operational. Finally, the managed MS must acquire its operational parameters with TFTP.

After a managed MS has obtained its operational parameters, or after an unmanaged MS has registered with the network, the MS preprovisioned service-flow connections are established.

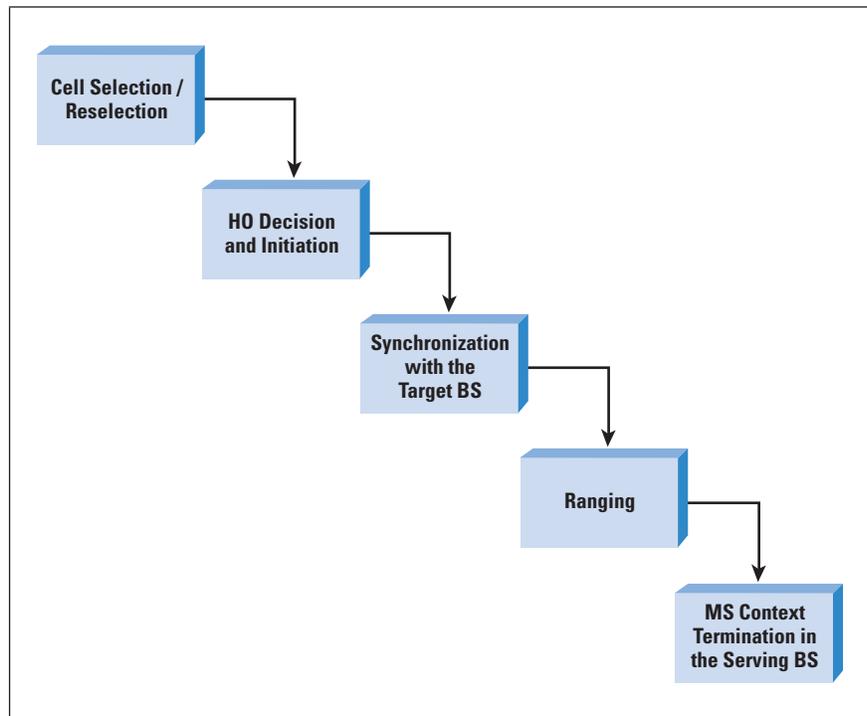
Mobility Support

As discussed previously, IEEE 802.16e introduced mobility support, defining an OFDMA PHY layer and signaling mechanisms to enable location and mobility management, paving the way for mobile WiMAX. The WiMAX Forum details four mobility scenarios in addition to the fixed WiMAX scenario. In the nomadic and portable mobility scenarios, the point of attachment of a fixed *Subscriber Station* (SS) can change. The simple mobility scenario allows MSs to roam within the coverage area with speeds up to 60 km/h, but handovers may cause connection interruptions of up to 1 second. In the so-called *full-mobility scenario*, the MS speed can be as much as 120 km/h, and transparent handovers are supported. This last scenario is what many might consider as the real mobile WiMAX scenario, but all five scenarios are “standards-compliant.”

Although three different types of handovers are defined in [2], *Hard Handover* (HHO), *Macro Diversity Handover* (MDHO), and *Fast Base Station Switching* (FBSS), only HHO is mandatory for all mobile WiMAX equipment. This type of handover is often referred to as a *break-before-make handover*: first, the MS disconnects from the serving BS and then connects to the target BS. Because of the short disconnection period, packets may be lost; HHO is less sophisticated than either MDHO or FBSS and may be inappropriate for some applications. The MS must also register with the target BS and reauthenticate with the network, typically meaning further delays before actual data exchange can (re)start. If multiple handover types are supported and enabled, the BS decides which type should take precedence over the other. MDHO and FBSS are enabled or disabled during the registration of the MS with the BS.

Figure 5 illustrates the five stages of a successful HHO in mobile WiMAX. The first stage is to select the target BS cell based on information about the network topology surrounding the serving BS through periodically broadcasted neighbor advertisements. The advertisements include the same information on the serving BS neighbors that the *Downlink Channel Descriptor* (DCD) and *Uplink Channel Descriptor* (UCD) messages of the neighboring BSs would include. For example, a neighbor advertisement message includes channel information of the neighboring BSs so that the MS can synchronize with them and perform scanning operations to evaluate their suitability as potential targets for a HO.

Figure 5: The Five Phases of a Successful HHO



The second phase is to make the actual decision to initiate the handover procedure, when a certain network (say, congestion in the serving cell requires load balancing) or channel condition threshold (for example, low received *Signal-to-Interference + Noise Ratio* [SINR] in the current cell) is crossed. The actual decision to start the message exchange for the MS to migrate from the radio interface of the serving BS to the radio interface of another BS can be made by the MS, BS, or the network. In the third phase, the MS synchronizes with the downlink transmission of the target BS and obtains the transmission parameters for the downlink and the uplink. The time consumed to perform the synchronization procedure depends on the amount of information the MS received about the target BS in the neighbor advertisement messages prior to the handover. The average synchronization latency without previously acquired information about the target BS ranges from two to three frame cycles, or approximately 4 to 40 ms depending on the OFDMA frame duration used in the system. The more extensive the channel parameter list received in the neighbor advertisement messages prior to the handover, the shorter the time to achieve the synchronization.

After synchronizing, the MS and the target BS initiate the ranging procedure. During this fourth step in HHO, MS and BS exchange the required information so that the MS can reenter the network. The target BS can request information about the MS from the (previously) serving BS and other network entities. Again, the more information made available to the target BS, the shorter the time to reenter the network, because the target BS may skip some steps from the network (re)entry procedure described earlier. In short, sharing context information before the actual handover optimizes the handover procedure and decreases its latency. In the last step of a HHO, the MS context at the serving BS is terminated and resources are released.

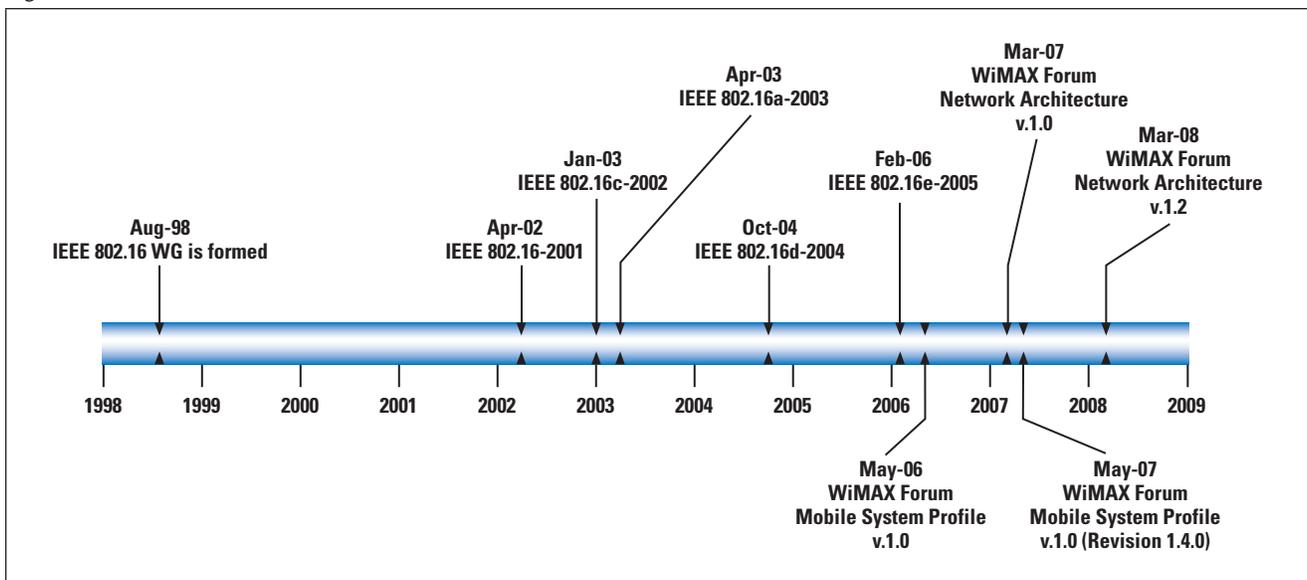
If MDHO and FBSS are supported, the following stages, in addition to those already described in the HHO procedure, must be performed: (a) decision to enable MDHO or FBSS, (b) diversity set update, and (c) anchor BS selection. In macrodiversity communications the MS maintains a connection to one or more serving BSs simultaneously, enabling soft or make-before-break handovers. In [2], the transition of the MS from the air interface of one or more serving BSs to the air interface of one or more target BSs is referred to as a MDHO. The MS and the BS both maintain a list called the *diversity set*, which includes all serving BSs involved in the MDHO communication. The MS maintains both uplink and downlink unicast connections to all the BSs in the diversity set, and one of the serving BSs is defined as the anchor BS. Note that all BSs involved in the diversity set use the same set of CIDs for the connections established between the MS and the serving BSs.

In FBSS, the MS transmits to and receives data from a single serving BS during any frame period. The BS, to which the MS has the connection to at any given frame, is called the *anchor BS*. The MS maintains a diversity set, which includes all active BSs in its range, and can change its anchor BS on a frame-by-frame basis, based on certain criteria. The transition from the serving anchor BS to the target anchor BS in FBSS is done without invocation of the normal handover procedure, and only the anchor BS update procedure is needed. After all, the MS has collected all required information about all BSs during the diversity set update ranging procedures.

Mobile WiMAX vs. HSPA

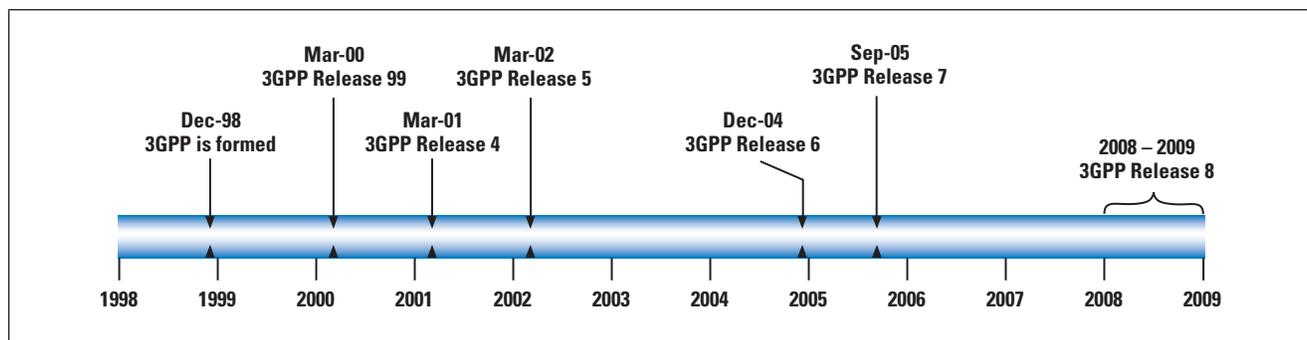
Mobile WiMAX and *High-Speed Packet Access* (HSPA) are expected to be the two major contestants in the rapidly growing wireless broadband market. The two, however, come from different origins. Figure 6 summarizes the evolution toward mobile WiMAX. It all started with the establishment in August 1998 of the IEEE 802.16 working group, which published its first standard (IEEE 802.16-2001) in April 2002. This first version defines a single carrier system operating in the 10- to 66-GHz frequency band and only under *line-of-sight* (LOS) conditions. The IEEE 802.16c-2002 amendment detailed system profiles for the original standard based on the 10- to 66-GHz frequency band. IEEE 802.16a-2003 introduced support for 2- to 11-GHz frequencies and *non-line of sight* (NLOS) operation, and adopted the use of OFDM and OFDMA. IEEE 802.16d-2004^[1] consolidated all these previous versions and amendments in a single document, and further enhanced the system. Fixed WiMAX is based on IEEE 802.16d-2004, [3], and [4]. Mobile WiMAX is based on the IEEE 802.16e-2005 amendment^[2], which introduced mobility support, as well on [3] and [4].

Figure 6: The Road Toward Mobile WiMAX



HSPA is a set of technological enhancements to the already widely deployed *Wideband Code Division Multiple Access* (WCDMA) cellular networks defined by the *Third Generation Partnership Project* (3GPP). Figure 7 illustrates the WCDMA specification evolution. The origins of HSPA can be traced in the foundation of 3GPP in December 1998. The original aim of 3GPP was to develop a third-generation WCDMA system, and in the process, HSPA was introduced. In March 2000, Release 99, the original standard specifying the WCDMA system, was published. A year later, the first enhancements were published in Release 4, which introduced, among others, an IP-based core network. Release 5 introduced *High-Speed Downlink Packet Access* (HSDPA) and defined the *3GPP IP Multimedia Subsystem* (IMS). *High-Speed Uplink Packet Access* (HSUPA) and some further improvements to HSDPA were defined in Release 6 (December 2004). Release 7 further enhanced QoS support and defined mechanisms to decrease network latency. Release 8 is expected to be published in 2008, and it will include specifications for the next step, called *3GPP Long-Term Evolution* (LTE). LTE is meant to deliver maximum cell throughputs an order of magnitude larger than HSPA.

Figure 7: The Evolution of the 3GPP WCDMA Standard



Mobile WiMAX evolved out of a broadband wireless LAN/MAN technology, and vendors currently report that it can deliver maximum cell capacities of 46 and 7 Mbps in downlink and uplink transmissions, respectively. However, mobility management is a later addition and, according to Maravedis, by September 2007 only 12 percent of all deployed *Customer Premises Equipment* (CPE) was IEEE 802.16e-2005-compliant^[16]. On the other hand, HSPA is based on a solid foundation of mobility management techniques with wide deployment in cellular networks around the globe, but can currently deliver maximum cell throughputs of only 14.4 and 5.8 Mbps in downlink and uplink transmissions, respectively.

Either commercial or trial networks of both technologies have already been implemented all over the world. However, according to the *Global Mobile Suppliers Association* (GSA), HSPA networks have yet to be deployed in China and India, both of which are large and rapidly growing market areas for wireless communications. According to Maravedis, both India and China have at least WiMAX trial deployments in place.

As mentioned already, the vast majority of current WiMAX deployments do not support mobility. Up to now, fixed WiMAX has been used mainly for last-mile broadband connectivity for sparsely populated rural areas. The largest commercial IEEE 802.16e-2005-compliant system is currently the *Wireless Broadband (WiBro)*^[17] network in South Korea, which supports simple mobility up to 60 km/h. Even though WiMAX and WiBro are both based on the same standards, WiBro was developed by the South Korean telecommunications industry before the WiMAX Forum adopted mobility support for its system profiles. WiMAX and WiBro are often cited as separate technologies, even though cooperation is in place in order to assure interoperability between the two.

Summary

In this article we presented an overview of mobile WiMAX, a much-heralded technology for next-generation mobile broadband networks; mobile WiMAX is an intricate system. We introduced WiMAX and the role of the WiMAX Forum, and summarized the important points of the WiMAX network reference model and the PHY and MAC layers. We addressed mobility support, but not the security aspects. Finally, we briefly compared WiMAX with HSPA, presenting their respective evolutions and illustrating their worldwide deployments. We hope that this article will serve as a valuable primer, and we highly recommend that those interested in the mobile WiMAX technology check the bibliography.

Bibliography

- [1] IEEE 802.16 Working Group, “IEEE Standard for Local and Metropolitan Area Networks, Part 16: Air Interface for Fixed Broadband Wireless Access Systems,” IEEE Standard 802.16-2004, October 2004.
- [2] IEEE 802.16 Working Group, “IEEE Standard for Local and Metropolitan Area Networks, Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems, Amendment 2: Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands,” IEEE Standard 802.16e-2005, February 2006.
- [3] WiMAX Forum Network Working Group, “WiMAX Forum Network Architecture—Stage 2: Architecture Tenets, Reference Model and Reference Points—Release 1, Version 1.2,” WiMAX Forum, January 2008.
- [4] WiMAX Forum Network Working Group, “WiMAX Forum Network Architecture—Stage 3: Detailed Protocols and Procedures—Release 1, Version 1.2,” WiMAX Forum, January 2008.

- [5] K. Pentikousis, “Wireless Data Networks,” *The Internet Protocol Journal*, Volume 8, No. 1, March 2005, pp. 6–14.
- [6] IEEE 802.16 Working Group, “IEEE Standard for Local and Metropolitan Area Networks. Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems. Amendment 3: Management Plane Procedures and Services,” IEEE Standard 802.16g-2007, December 2007.
- [7] C. Perkins (Ed.), “IP Mobility Support for IPv4,” RFC 3344, August 2002.
- [8] D. Johnson, C. Perkins, and J. Arkko, “Mobility Support in IPv6,” RFC 3775, June 2004.
- [9] K. Leung, G. Domemety, P. Yegani, and K. Chowdhury, “WiMAX Forum/3GPP2 Proxy Mobile IPv4,” Internet-Draft, Work in Progress.
- [10] R. Droms, “Dynamic Host Configuration Protocol,” RFC 2131, March 1997.
- [11] R. Droms (Ed.), J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney “Dynamic Host Configuration Protocol for IPv6 (DHCPv6),” RFC 3315, July 2003.
- [12] S. Thomson and T. Narten, “IPv6 Stateless Address Autoconfiguration,” RFC 2462, December 1998.
- [13] K. Sollins, “The TFTP Protocol (Revision 2),” RFC 1350, July 1992.
- [14] J. Case, M. Fedor, M. Schoffstall, and J. Davin “A Simple Network Management Protocol (SNMP),” RFC 1157, May 1990.
- [15] J. Postel and K. Harrenstien, “Time Protocol,” RFC 868, May 1983.
- [16] K. Pentikousis, J. Pinola, E. Piri, F. Fitzek, T. Nissilä, and I. Harjula, “Empirical Evaluation of VoIP Aggregation over a Fixed WiMAX Testbed,” Proceedings of The 4th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (TRIDENTCOM), 18–20 March, 2008, Innsbruck, Austria.
- [17] Telecommunications Technology Association, “Specifications for 2.3GHz Band Portable Internet (WiBro™) Service,” TTA Standard TTAS.KO-06.0082/R1, December 2005.

JARNO PINOLA received his M.Sc. from the University of Oulu, Oulu, Finland, in Spring 2008. During his studies, he specialized in telecommunication systems and wrote his Master's Thesis on mobility management issues in wireless broadband systems. Currently he is working as a Research Scientist at VTT Technical Research Centre of Finland in Oulu, Finland. He can be contacted via e-mail at: [**jarno.pinola@vtt.fi**](mailto:jarno.pinola@vtt.fi)

KOSTAS PENTIKOUSIS studied computer science at Aristotle University of Thessaloniki, Thessaloniki, Greece (B.Sc. 1996), and Stony Brook University, Stony Brook, New York, USA (M.Sc. 2000, Ph.D. 2004). He is a tenured Senior Research Scientist at VTT Technical Research Centre of Finland, in Oulu, Finland. He has published internationally in several areas, including mobile computing (mobility triggers, multiaccess, media-independent handovers, and energy consumption); transport protocols; applications; network traffic measurements and analysis; and simulation and modeling. Visit [**http://ipv6.willab.fi/kostas**](http://ipv6.willab.fi/kostas) for more information and contact details.

Letters to the Editor

IDNs

The DNS protocol is 8-bit clean (“Internationalizing the Domain Name System,” IPJ, Volume 11, No. 1, March 2008), even if some DNS clients and servers are not. The hardest thing about changing any Internet protocol is coordinating clients and servers during the transition.

And yet, with the DNS, no transition is needed to support UTF-8 domain names. If you want to publish a UTF-8 domain name, then run a name server that supports UTF-8. If you want to be able to access domain names in your own language, switch to DNS software that supports it. Implementations that are 8-bit clean are already available; ordinary market mechanisms will handle the rest.

Punycode is a gross hack that makes my stomach roil. You know it, I know it, any engineer will agree with you, so how did it get through the IETF?

The argument for where to stop internationalization does not spread to `protocol://` because it’s “gobble-de-gook” in English, too. Dots are a completely arbitrary character used to separate the hierarchy. There’s plenty of space at the top for UTF-8 names.

The real problem with IDN is homoglyphs.

—Russ Nelson,
nelson@crynwr.com

The author responds:

It would certainly make more sense in terms of design elegance and minimalism within the DNS if the label that was stored in the DNS was precisely the same label that was used in the interface between applications and the DNS client software. There is something rather clumsy about the approach that stores an encoded version of a canonical version of the label value, and relies on the application being capable of performing the *stringprep* and encoding functions in consistent and uniform ways. The resultant limitations on what can actually sit in DNS labels on a language-by-language basis are, in part, an outcome of the potential indeterminism of this canonicalization function.

But indeterminism is not a tolerable outcome of the DNS. The DNS is not a guessing game, and inconsistencies in the mapped transforms that are provided by the DNS trigger intolerable insecurities in the networked environment. So the *nameprep* profiles and the related restrictions on allowable Unicode code points are unavoidable if we want to avoid this indeterminism in the DNS.

So if *nameprep* is required in any case, then what we are left with to consider is the decision to use the Punycode *ASCII Compatible Encoding* (ACE) to map Unicode labels into the *Letter-Digit-Hyphen* (LDH) subset of ASCII. But is the Punycode ACE really that much of a problem? Within the overall IDN framework the Punycode algorithm is not so complex that the risk of incorrect implementations is significant, the algorithm is not processor-intensive, and the outcome does not inflate the encoded labels to an impossible length. The advantage of Punycode is that the DNS servers do not require modification, and the clients that manipulate IDNs required additional *nameprep* functions in any case, so Punycode was evidently intended to be the least-impact approach that spared DNS servers from a potential requirement for modification.

To me, this solution appears to be a design tradeoff, in so far as the ACE approach circumvents the observed problem of non-8-bit clean DNS servers sitting within the deployed DNS, and does not in and of itself demand novel roles and functions on the part of the clients of the DNS in addition to what was already necessitated by the IDN *nameprep* function. However, at the same time it creates an annoying inconsistency in the overall framework of the design of the DNS, where certain labels in the DNS are intended to trigger a Punycode transform into an equivalent Unicode string while other labels are meant to be used without further transforms applied.

My judgment of the short-term path of least risk sits with the ACE approach as adopted for IDNs, but at the same time I agree with Russ' discomfort that the path that preserves the long-term essential broad utility and function of the DNS through consistency of design and application sits in an 8-bit clean DNS without the adornment of any form of an ACE.

And, yes, I agree with Russ that the most significant problem with IDNs is homoglyphs, because of continued reliance of an underlying approach of “appearance is everything” in terms of the integrity of the DNS as an identity framework.

—Geoff Huston,
gih@apnic.net

More IDNs

The LDH restriction referred to in “Internationalizing the Domain Name System” (IPJ, Volume 11, No. 1, March 2008) was relaxed in RFC 1123^[1] to allow a host name to begin with either a letter or a digit.

—Andrew Friedman

[1] R. Braden, Editor, “Requirements for Internet Hosts – Application and Support,” RFC 1123, October 1989.

The author responds:

My thanks to Andrew for pointing this out. It has been commonly recounted that this relaxation of the LDH convention was associated with the successful registration of the DNS name `3com.com` and that the RFC paperwork was revised following this registration. Since then the most visible set of names that used this “liberal” revision of LDH with names that have leading digits were telephone number mapping name sets, including the venerable `tpc.int` domain of the early 1990s and, more recently, ENUM. As for names with leading hyphens, I don’t believe that we are at the point of allowing Morse code into the DNS yet, but I’m sure that someone somewhere is working on it!

—Geoff
(-- . . --- .-. .-.)

We want to hear from You

Your feedback is important to us. Please send your comments and suggestions to ipj@cisco.com. And don’t forget to visit our Website at <http://www.cisco.com/ipj> where you can read or download back issues, update and renew your subscription, and find articles using our index files. We also encourage you to participate in our online forum at <http://ipjforum.org>

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

Fragments

OLSR stands for *Optimized Link State Routing Protocol*.

DUMBO

The *Digital Ubiquitous Mobile Broadband OLSR* (DUMBO) project deploys mobile wireless networks on an ad hoc basis for emergency conditions, such as after a natural disaster when a fixed network infrastructure is not available.

A *Mobile ad hoc Network* (MANET) consists of mobile nodes that automatically cooperate to support the exchange of information through wireless medium. Since the MANET does not rely on fixed telecommunication infrastructure, it is suitable for emergency situations and can be set up in a short amount of time. Using lightweight portable mobile nodes, MANET coverage can penetrate deep into areas not easily accessible by roads or into areas where the telecommunication infrastructure has been destroyed.

DUMBO allows streaming video, *Voice over IP* (VoIP) and short messages to be simultaneously transmitted from a number of mobile laptops to a central command center, or to the other rescuers at the same or different disaster sites. The DUMBO command center has a face recognition module that identifies potential matches between unknown victims' face photos taken from the field and a collection of stored known face images. In addition, sensors can be deployed to measure environmental data such as temperature and humidity. Data from the sensors can be sent to the command center which analyzes or passes it on to the other mobile nodes. The command center can be located either in the disaster area or anywhere with Internet access. DUMBO technology is currently being deployed in cyclone-ravaged Burma. See <http://www.interlab.ait.ac.th/dumbo/> and <http://www.relief.asia/>

Upcoming Events

The *Internet Engineering Task Force* (IETF) will meet in Dublin, Ireland, July 27 – August 1 and in Minneapolis, Minnesota, November 16 – 21, see <http://www.ietf.org/>

APNIC, the *Asia Pacific Network Information Centre*, will hold its Open Policy meeting in Christchurch, New Zealand, August 25 – 29, see <http://www.apnic.net/meetings/26/>

[Ed.: I will be organizing a pipe organ demonstration event on August 26 as part of the opening reception for APNIC 26, see <http://organdemo.info>]

The *North American Network Operators' Group* (NANOG) will meet in Los Angeles, California, October 12 – 14. Immediately following the NANOG meeting, the *American Registry for Internet Numbers* (ARIN) will meet in the same location, October 15 – 17. See <http://nanog.org> and <http://arin.net>

The *Internet Corporation for Assigned Names and Numbers* (ICANN) will meet in Paris, France, June 22 – 26, and in Cairo, Egypt, November 2 – 7. See <http://icann.org>

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
Distinguished Career Professor of Computer Science and Public Policy
Carnegie Mellon University, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, General Chair Person, WIDE Project
Vice-President, Keio University
Professor, Faculty of Environmental Information
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

*The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc. www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com*

Copyright © 2008 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners.

Printed in the USA on recycled paper.



The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD U.S. Postage PAID PERMIT No. 5187 SAN JOSE, CA
--