Volume 07 Issue 03 Published, August 19, 2003 ISSN 1535-864X

Intel® Technology Journal

Wireless Technologies

Wireless technologies offer significant opportunity for the computing and communication industries. This issue of Intel Technology Journal (Volume 7, Issue 3) offers an in-depth look into Intel's innovation and work to support and enhance the new services made possible by Intel® Centrino™ mobile technology, Wireless LANs, and network co-existence.

Inside you'll find the following papers:

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Volume 07 Issue 03 Published, August 19, 2003 ISSN 1535-864X

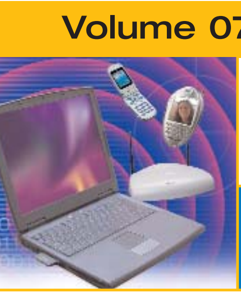

Intel® Technology Journal

Wireless Technologies

Articles

Preface: Wireless Technologies

Lin Chao
Publisher, Intel Technology Journal

In Aug. 2003, the planet Mars will be closer than ever in recorded history. Mars has begun an orbit with Earth that will, over the next year, provide the best viewing opportunity since Neanderthal pre-historic man looked skyward. On Aug. 27, 2003, at 5:51 AM (Eastern Daylight Time in USA), Mars will be 34.6 million miles (55,758,000 kilometers) from Earth. This will be the closest that Mars has come to our planet in nearly 60,000 years. Mars will shine some 85 times brighter and appear six times larger. The next time it will be this close will be Aug. 28, 2287. This is an opportunity of a lifetime for us to get an unprecedented look at Mars.

In technology, we are experiencing a close encounter with wireless technology by providing an order of magnitude–"10X"–improvement in ease of connectivity. Never before has wireless technology been so easy and so available as we connect at Starbucks and McDonald's. This issue of Intel Technology Journal (Volume 7, Issue 3) offers an in-depth look into Intel's innovation and work to support and enhance the new services made possible by Intel® Centrino™ mobile technology, Wireless LANs, and network co-existence.

The first paper, an invited paper, touches on "Open Innovation"–a model Intel often uses–which incorporates both external and internal ideas for product innovation. The next three papers delve into hotspots, co-existence and roaming. The second paper looks at hotspots and an open hotspot architecture based on open standards. The third paper explores how Wireless Personal Area Network (WPAN) technology, such as Ultra Wideband (UWB) or Bluetooth* technology, enables new usage models when used in conjunction with Wireless Local Area Networks (WLANs) technology. The fourth paper discusses how Dynamic Networks on Demand (DND) framework automates network creation for wired and wireless network topologies and executes resource transitions without ever plugging or unplugging a wire or configuring a network device.

The fifth paper addresses major candidate technologies for unlicensed band Wireless Local Area Network (WLAN) and digital and radio frequency technology improvements that will contribute to the commercialization of high-throughput wireless LAN systems, while the sixth paper looks at some of the challenges associated with a client solution and how quick and seamless connectivity to WLANs, including features to enhance the user experience, could be supported within the Intel® Wireless PROSet software framework.

Foreword: Wireless Technologies

Gadi Singer
Vice President, Wireless Communications and Computing Group
General Manager, PCA Components Group

Wireless has arrived, and with it the freedom and flexibility to connect us at work, in a hotel, at an airport, a restaurant, even on a train or plane. Wireless technology is shaping new realities and new behaviors as more and more people discover its benefits. Today, users demand instantaneous access to content anywhere, anytime. This usage model is driving an explosive demand in wireless-capable devices. In the cellular market alone, analysts predict over 200 million 2.5G and 3G wireless data-phone subscribers by the end of this year and over one billion subscribers within five years[1]. The 802.11 Wi-Fi hotspots have shown impressive growth and are quickly becoming the popular mechanism to access the Internet. Recently launched Intel® Centrino™ mobile technology with built-in 802.11b is experiencing great success and is being met with much excitement in the market place. All these factors clearly point toward wireless adoption at a mass-market level. This "great wireless wave" offers significant opportunity for the computing and communication industries. However, there are several challenges that the industry must overcome to make these technologies truly useful, fun, and seamless to use. This publication offers a peek into Intel's innovation and work to improve the user experience in various aspects of wireless technologies.

To meet specific usage model requirements, various wireless standards have evolved over the years. For example, initial cellular standards such as Global System for Mobile telecommunication (GSM) were used primarily for making voice calls. However, later technologies such as General Packet Radio Service (GPRS), EGPR, and Universal Mobile Telecommunications Service (UMTS) are adding the capabilities for high-rate data communication. 802.11 Wi-Fi has gained wide acceptance as the wireless WLAN standard of choice for laptop computers. Bluetooth* is emerging as the preferred standard for the Wireless Personal Area Network (WPAN). Convergence of these standards into coexistence within a single product is inevitable. Intel is developing the solutions to overcome the technological challenges with the vision of providing seamless connectivity to the users across multiple wireless standards on an open platform.

To ensure user satisfaction and to minimize Total Cost of Ownership (TCO), manageability of the wireless network is another area that needs focus. As wireless networks become ubiquitous, quick and easy connectivity to these networks is needed. Additionally, the setup of the wireless network must be simplified for mass adoption—especially in the home environment. One of the papers in this issue of Intel Technology Journal describes how Intel's Wireless PROSet software (which is a key component of the Intel® Centrino™ mobile technology) supports quick and seamless connectivity to WLANs, including additional features that enhance the user experience.

It is clear that wireless technologies offer significant opportunities for innovation, which in turn drive new capabilities. To ensure that innovation is unconstrained and to overcome various challenges, the industry must work toward developing open standards. Open standards with well-defined interfaces maximize innovation pace and remove interdependencies. Vendors are free to focus on areas where they can add value. Of course, to ensure the success of the overall platform based on open standards, the industry must create and adopt a strong interoperability infrastructure.

Some of the challenges such as security will require organized industry-wide effort and cooperation. For e-commerce to truly take hold in the wireless space, device security must be improved. Security has multiple dimensions that may require a phased approach. The initial phase could make the platform itself secure from external viruses and harmful content; a second phase could encompass secure communication; and the final phase could involve making the entire environment secure.

We are in the initial stages of a wireless growth spiral characterized by user desire to access content anywhere, anytime. This ongoing wireless adoption will require new infrastructure build-out. New business models and life styles, which we can only imagine today, will emerge. Of course, industry will need to overcome challenges as we continue on this exciting journey. Intel has played a key leadership role in the emerging wireless landscape and will continue to do so, both in shaping the new world and in delivering best-in-class technologies, capabilities and products.

The great wireless wave is here; let's ride it high!

* Other names and brands may be claimed as the property of others.

®Intel and Centrino are registered trademarks and trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

[1] Source: IDC, Gartner and ARP Group

Open Platform Innovation:
Creating Value from Internal and External Innovation

Henry Chesbrough, Haas School of Business, University of California - Berkeley

Index words: innovation, management, metrics, open innovation

ABSTRACT

This paper explains the Open Innovation approach to managing innovation, with particular emphasis on the use of platform technologies in the Information Technology (IT) sector and, more specifically, on the personal computer industry. Platforms provide an architecture to combine internal and external innovations in ways that create value throughout the chain of activities that deliver a useful technology to the market. Value creation is vital for adoption of the architecture by third parties and customers. The architecture must also enable the platform architect to capture a portion of the value created. This value capture is critical to sustaining the advance of the platform.

The value of building a platform technology will be illustrated briefly by the early history of Adobe Systems, while the difficulty of doing so will be demonstrated by a comparison with another early software company, Metaphor Computer Systems. Implications for other IT platform architects, such as Intel, will be examined, including the appropriate metrics for managing an Open Innovation approach.

INTRODUCTION

The process of industrial innovation has undergone a significant shift, since the heyday of Vannevar Bush's vision of "Science–The Endless Frontier" [1]. That vision assumed that industrial corporations would be able to innovate only by conducting basic research activities internally, and by carrying the results of that research through to the market.

While that approach worked well for a variety of US corporations during the 20th century, as demonstrated by business historian Alred Chandler [2], events in the past thirty years have eroded the conditions that supported that approach. As explained in more detail elsewhere [3], these conditions include

- The increasing mobility of technical and managerial personnel across firms.
- The rising quality and relevance of university research.
- The explosion in college graduates and the increasing quality and quantity of human capital.
- The growth in the quality and quantity of international research.
- The dramatic growth in venture capital and private equity, enabling startup companies to attract high-quality talent.

These erosion factors have rendered the internally focused model of Closed Innovation obsolete in most industries. Today, the earlier model has been overtaken by a model of Open Innovation. Open Innovation can be briefly defined as utilizing external as well as internal ideas as inputs to the innovation process, combined with employing internal and external paths to market for the results of innovative activities.

Figure 1 depicts this new Open Innovation paradigm. It shows internal and external ideas flowing into the R&D process, and it shows the outputs of that process going to market through external paths, in addition to the internal path. Of particular importance is the flow of ideas and technologies into and out of the process throughout. Ideas can come into the process, for example, from internal research investigations, from external research, from licensing in another company's technology, or from an acquisition of a company's product. Similarly, ideas can flow out of the process to market in numerous ways. Many go to market through the company's own channels, while others may be licensed out, or spun out into a new venture, or into a new joint venture.

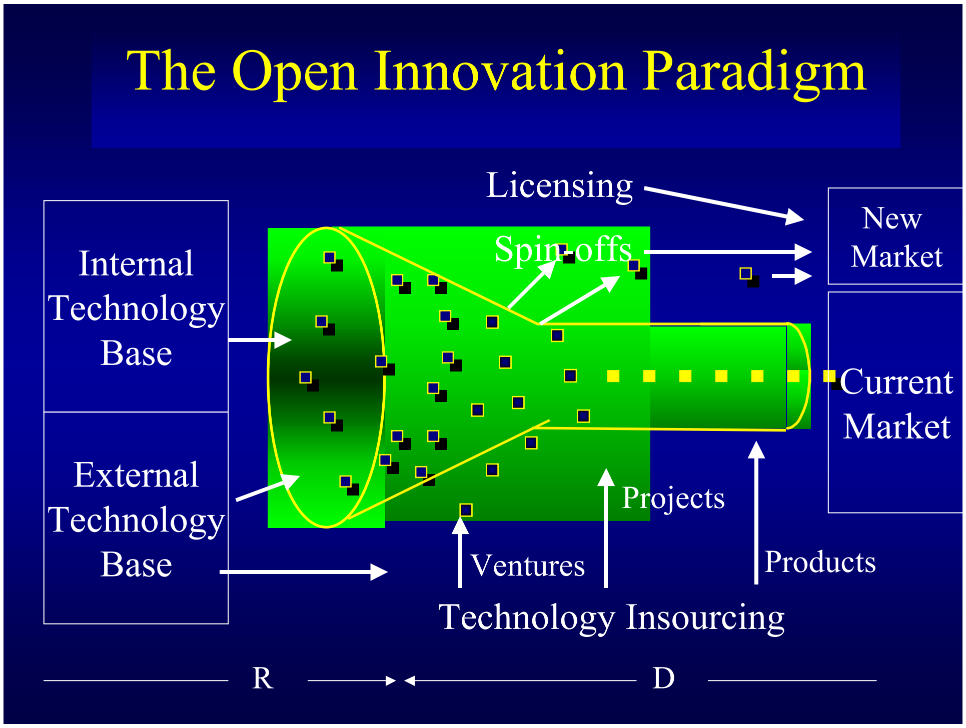

Figure 1: Open innovation paradigm

The opportunities and challenges of managing Open Innovation processes have particular relevance to the IT sector of the economy. The example of Platform Technologies [4] is used to illustrate the value and the difficulty of this approach to managing industrial innovation. The contrasting experiences of Adobe and Metaphor will explicate these issues, and the paper closes with some observations regarding the new and different metrics needed to manage a platform technology effectively.

PLATFORM TECHNOLOGIES AND THE BUSINESS MODEL

A technology by itself is worth very little. Most patents are worth almost nothing, while a very select few are worth an enormous amount [5]. While some of this disparity may be due to the difference in the quality of the technology, business history is full of examples where an inferior technology overtook a superior technology in a battle to set a standard [6].

A significant portion of the success of an inferior technology over a superior one results from differences in the business model employed to commercialize the two technologies. While a detailed definition of a business model can be found elsewhere [7], its key elements for the purposes of this paper are 1) the creation of value through a chain of activities that extend from raw materials through to the end customer of the product or service; and 2) the capture of a portion of that value within the chain of activities.

The first of these two requirements requires an architecture that combines a variety of constituent parts in a coherent way to form a complex system. In any such system, the number of possible combinations of these parts vastly exceeds an organization's ability to try them all. Heuristics and experience are employed to restrict the number of feasible solutions evaluated. "Heuristics" means human and organizational cognition plays an important role in the creation of an architecture. Experience implies that the prior history of a company and an industry play a continuing role in shaping the types of future architectures that are created within an industry.

Today, we have the tremendous business success of companies such as Intel and Microsoft in creating and managing the many combinations of Intel® X86 and Pentium® microprocessors with Microsoft's DOS* and

®Intel and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other brands and names are the property of their respective owners.

Windows[*] operating systems. The success of this platform has drawn attention to the value of an effective, widely supported platform. What is perhaps forgotten is how tremendously difficult it was to establish a platform in the first place, and then how difficult it is to sustain the advance of that platform over time.

We review a brief history of two software companies, Adobe Systems and Metaphor Computer Systems, and use the comparative history to explore the difficulty of creating and sustaining a platform over time. Those interested in a more extensive treatment of these companies can find it elsewhere [7].

Adobe Systems' Early History

Today Adobe Systems is known as one of the most profitable and most valuable (as measured by its stock market value) software companies. The company's first important product, PostScript, created a new industry segment within the personal computer industry, known as *desktop publishing*. This has been an important platform technology since the mid-1980s.

In its infancy, however, its technology was developed within Xerox's Palo Alto Research Center (PARC). The journey from internal laboratory to successful software giant was far from straightforward and illustrates the challenges of creating a platform technology.

Within PARC, what was then called Interpress was a means of allowing Xerox's laser printers to print what was displayed on a Xerox Star workstation. This capability became known as WYSIWYG, or "what you see is what you get." Thus, the Interpress technology was a component of the entire Xerox system, consisting of a number of networked workstations, connected to a shared laser printer via what would become known as the Ethernet networking protocol.

John Warnock and Charles Geschke, who both worked on this technology while at PARC, wanted to create a standard around Interpress. However, their management within Xerox resisted this, because they did not want to give away one of the primary differentiating features of the Star system. Eventually, Warnock and Geschke would leave PARC over this dispute.

Once outside of Xerox, the two had to create a business to exploit the technology that they labeled PostScript*. Initially, they decided to create a turnkey publishing system, complete with proprietary hardware, software, fonts, and applications. This complete systems approach was very similar to that of Xerox in its Star workstation. As they tried to raise money to fund their venture, though, Warnock and Gescke ran into resistance from venture backers, as well as from leading companies of the day. Both Gordon Bell of Digital Equipment Corporation and Steve Jobs of Apple Computer persuaded them to pare down the scope of their technology efforts and to focus upon the fonts themselves. Other companies would work with Adobe to utilize the PostScript fonts in their products. Eventually, Hewlett Packard and Canon agreed to bundle in PostScript as a standard item in their laser printers, while Apple agreed to support PostScript in its software. The PostScript platform was now in place, and Adobe grew to become a valuable company. So valuable, in fact, that Microsoft would later team with Apple to create a rival standard known as TrueType*, to take away some of the profits from the PostScript platform.

Metaphor's Illustrative Failure

Metaphor Computer is another software company that was started by people who were previously employed at Xerox. The two founders were Donald Massaro and David Liddle, both of whom were highly regarded in their roles at Xerox. They had envisioned a user-friendly software approach to construct database queries, to allow analysts to address a number of questions that were currently being held up by the long queue of projects within the MIS organization. It was to become one of the very first client-server applications.

Like Adobe, Metaphor initially chose to construct the entire computing system environment to support their user-friendly query construction approach. And Metaphor persisted for a number of years in pursuing a "do it all yourself" approach to implementing their technology.

Unlike Adobe, Metaphor chose not to partner with other firms to commercialize its technology. If a customer valued the ability to generate database queries of corporate mainframe data, the customer had to buy the entire stack of Metaphor's offering. There was no layer that was carved out and shared with other firms. The lack of third-party support greatly restricted Metaphor's market penetration, and the company was sold in 1991. The Metaphor product was withdrawn from the market.

Challenges in Building a Platform Business Model

Why did Adobe and Metaphor initially choose to implement entire systems? This can be explained in part by the Xerox PARC legacy, where there were researchers working on basic physics and materials, as well as computer scientists, and even anthropologists and sociologists. PARC celebrated the multi-disciplinary scope of its research capability, and successful

[*] Other brands and names are the property of their respective owners.

researchers from that institution may have carried that value forward with them. And Xerox, its parent company, ran its business by selling complete solutions to its (primarily) corporate customers. This entailed pricing at gross margins in excess of 50%, maintaining a direct sales force, and performing all of the R&D necessary to advance the product line.

But a more general reason is that creating new architectures to connect disparate parts is a complex and interdependent undertaking. Initially, the best way to partition the functions within a system is far from clear, and changes in one part of the system often create problems in other areas, even those that might seem at first to have no connection to the change. It helps greatly to be able to manage this interdependency inside a single firm. The many tradeoffs and compromises that must be made do not involve any monetary payoffs to outside organizations.

Why, then, did Adobe eventually shift away from a turnkey approach, and build a platform instead? Adobe was able to mature its technology sufficiently to envision partitions and interfaces between the constituent elements of its offering. It readily appreciated the greater market penetration it would realize if Canon, HP, and Apple embedded the PostScript library in their printers. It also focused its business model to deliver and enhance the fonts themselves, leaving the other functions to others in the platform. This allowed it to make money, without having to charge too high a price to cover unnecessary overheads that weren't directly related to its value added.

Once a platform is established, new challenges arise. On the one hand, the platform architect wants to encourage as many third-party developers as possible to make investments to create products that leverage the platform, thus increasing its value. This leads to a policy of neutrality towards any specific developer.

On the other hand, the architect cannot order third parties to develop any specific applications, or to make improvements to existing applications. If important applications are missing or inferior, the platform could be at risk of being supplanted by a rival platform. This leads to a policy of favoritism towards certain developers, and even selective forward integration into making the application. If the desired new or improved application requires substantial interdependency with the architecture of the platform itself, this again encourages closer organizational integration to execute the task.

METRICS FOR MANAGING OPEN PLATFORM INNOVATION

It is a well-known axiom of management that, in order to get better, an activity must first be measured before it can be improved. (W. Edwards Deming's work in quality improvement is but one example of this.) How then, can one measure open platform innovation, in order to manage it more effectively?

The first step is to realize that many of the traditional measures of R&D exclude vital elements of a healthy, Open Innovation process. Consider the percentage of sales spent on R&D. A moment's reflection reveals that this only captures the spending on *internal* R&D and omits entirely any investment by any other party, including suppliers, customers, or third parties. Yet the amount of external spending may be a critical leading indicator of the vitality of one's platform. Few, if any, companies today track this number, and it is ignored by Wall Street analysts, but it is as important to the value of one's platform as the amount of internal money one spends. Indeed, outside financial analysts would do well to track such external spending in their own evaluation of the vitality of a company's technology, and by inference, that company's stock price.

Another obsolete metric is the number of patents generated from internal R&D. While patents can be useful and a few of them can turn out to be quite valuable, becoming a "patent mill" is a poor way to manage industrial innovation. It is far better to track how many patented technologies are used in one's own products—as well as in other companies' products. It is after all the discipline of getting products to market that is vital to finishing the innovation process and delivering value to the market.

Furthermore, many companies pay small awards to researchers when those researchers are awarded a patent for work they have done at the company. However, there is usually no equivalent recognition given to an employee who finds an equally useful external technology. A bounty program for locating and accessing such technology makes good sense.

A few companies have gone still further, tracking the extent to which their portfolio of upcoming R&D projects utilizes external ideas, as well as ones generated internally. One company discovered that its internal ideas were 90% of the portfolio, and it decided to raise the external portion of the pipeline to 50% over the next five years. Other companies are setting stretch goals to increase the number of new products coming to market from R&D each year, which has the effect of motivating the R&D organization to supplement its internal R&D projects with additional ones located on the outside.

CONCLUSION

Technology companies create substantial value when they are able to establish platform technologies that are broadly accepted. However, these platforms are challenging to create and also challenging to advance once they are established. Companies will need to blend some amount of internal organization with open coordination with outside companies to manage these challenges. In an innovation environment that is characterized by broadly distributed knowledge, industrial companies will need to open up their innovation processes to leverage the wealth of external knowledge available. In the process, new measures and metrics will be required as well.

ACKNOWLEDGMENTS

I acknowledge the encouragement of David Tennenhouse of Intel in writing this article, and I thank Lin Chao for her patient editorial assistance.

REFERENCES

[1] Bush, Vannevar, "Science−The Endless Frontier," US Government Printing Office, July 1945.

[2] Chandler, Alfred, *Scale and Scope*, Harvard University Press: Cambridge, MA, 1962.

[3]Chesbrough, Henry, *Open Innovation: The New Imperative for Creating and Profiting from Technology*, Harvard Business School Press: Boston, MA, 2003.

[4] Gawer, Annabelle, and Cusamano, Michael, *Platform Technologies: How Intel, Microsoft and Cisco Drive Industry Innovation*, Harvard Business School Press: Boston, MA, 2002.

[5] Harhoff, D., and Scherer, F. M., "Technology Policy for a World of Skew-Distributed Outcomes," *Research Policy*, pp. 559-566, April 2000.

[6] Cusumano, M., Y. Mylonadis, and R. S. Rosenbloom, "Strategic Maneuvering and Mass Marketing Dynamics: The Triumph of VHS over Beta," *Business History Review* 66, no. 1, Spring 1992.

[7] Chesbrough, Henry, and Rosenbloom, Richard, "The Role of the Business Model in Capturing Value from Innovation: Evidence from Xerox Corporation's Technology Spinoff Companies," *Industrial and Corporate Change,* vol. 11 (3): pp. 529-555, 2002.

AUTHOR'S BIOGRAPHY

Henry Chesbrough teaches the management of innovation and technology at the Haas School of Business at the University of California – Berkeley. Previously, he was an assistant professor at the Harvard Business School. He is the author of *Open Innovation: The New Imperative for Creating and Profiting from Technology* (Harvard Business School Press, 2003). His e-mail address is chesbrou@haas.berkeley.edu

Public WLAN Hotspot Deployment and Interworking

Prakash Iyer, Corporate Technology Group, Intel Corporation
Victor Lortz, Corporate Technology Group, Intel Corporation
Lee Tapper, Corporate Technology Group, Intel Corporation
Roger Chandler, Corporate Technology Group, Intel Corporation
Roxanne Gryder, Corporate Technology Group, Intel Corporation

Index words: WLAN, 802.11, Wi-Fi, Roaming, Hotspot, WPA

ABSTRACT

Wireless data networking is becoming more and more popular, and network operators of various kinds (cellular providers, wireless Internet Service Providers, etc.) are beginning to deploy public Wireless Local Area Network (WLAN) hotspots around the world. However, broad adoption of these hotspots may be inhibited by technical obstacles such as ease of use, security, and the inability of users to roam across hotspots as they can with cell phones today. The latter problem in particular highlights the need for a common hotspot architecture that is based on open standards and is acceptable to different service provider communities. Such an architecture must also be flexible to accommodate users with a variety of mobile device form factors and login credential types, as well as different billing models.

We begin with a brief survey of the state and deployment of hotspots today and go on to describe a unified public hotspot architecture that addresses the technical obstacles mentioned above. A major theme of the paper is the transition from the insecure Universal Access Method (UAM) to more robust authentication and link security based on Wi-Fi[*] Protected Access (WPA*). While much of the discussion centers on authentication and authorization for internet access, we also touch upon issues that need to be addressed to enable more advanced services to be deployed and accessed from such hotspots in the future.

INTRODUCTION

Deployment of public Wireless Local Area Network (WLAN) hotspots, initiated by a diverse set of incumbent operators—cellular carriers (GSM and CDMA), Wireless Internet Service Providers (WISP), dial-up aggregators, and fixed broadband operators (xDSL, cable)—is growing rapidly across the globe. The predicted rate of deployment of WLAN technologies is impressive. The analyst firm Gartner predicts that by the year 2008 there will be more than 167 thousand public WLAN hotspots around the globe. In addition, there will be over 75 million users of public WLAN hotspots worldwide [1].

While the outlook for WLANs appears to be promising, there are several factors that may limit their viability as effective global solutions for wireless data connectivity. The following observations are worth noting:

- Each operator/carrier community has its own business models and independent standards' forums that are enabling "WLAN roaming and interworking" scoped primarily for that community.
- Hotspot deployment in urban areas is unlikely to be monopolized by individual operators or operator communities—limiting the available footprint for users—unless a common roaming framework is deployed. Therefore, intra-city roaming for WLAN users will be required if providers are to expand the use of their hotspots. Moreover, hotspot deployment has great potential for revenue generation, *a la* roaming in the cellular world.
- The smaller cell sizes, the low cost of equipment, and the lack of regulatory barriers for WLAN deployment encourage a greater diversity of operators to enter the business. Consequently, roaming will likely become more common for public WLAN users than for cellular users.
- With technology evolving rapidly, there is a substantial risk of fragmentation from this early

[*] Other brands and names are the property of their respective owners.

deployment of hotspots, thus inhibiting regional and global interoperability.

- User adoption may be slow if public WLAN services are not cost effective, widely available, secure, and easy to use.

These observations highlight the need to establish a common hotspot architecture that all operator types can embrace. In the remainder of this paper we describe a conceptual blueprint for hotspots, discuss authentication and security issues, and consider billing and settlement architectures to enable worldwide one-bill roaming for public WLAN.

THE PROPOSED HOTSPOT BLUEPRINT

Figure 1 is a conceptual illustration of a common hotspot blueprint that shows how a single hotspot could support roaming users with accounts managed by a wide variety of home operator types. For this approach to be practical, the authentication mechanisms and Authentication, Authorization, and Accounting (AAA) signaling between the hotspot and the different back-end authentication systems of different operator types must be compatible.

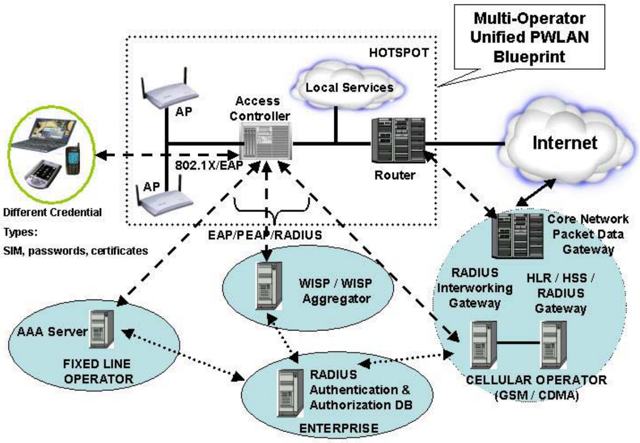

Figure 1: Conceptual hotspot blueprint

If a consistent approach for hotspot authentication, accounting, and billing can be established in the industry, it will become much easier for home operators of various types to begin offering public WLAN service to their customers. Home providers do not need to provide wireless network services in this model. All the home provider needs to do is to deploy AAA and billing infrastructure and to establish roaming agreements with one or more WLAN operators.

The State of WLAN Deployment Today

While the global proliferation of Wireless Local Area Networks (WLANs) continues at a rapid pace, the methods by which these networks are deployed—particularly in the public domain—are diverse and somewhat chaotic. According to the Wi-Fi Alliance[1], the most prevalent form of access today is based on web-browser hijacking and is referred to as the Universal Access Method (UAM). With browser hijacking, the hotspot redirects the user's browser to a local web server secured by TLS [2] (the standard security mechanism for web pages). The user's identity is authenticated to the UAM login page by entering a username and password on a form sent to the web server. Significant advantages of this method are ease of deployment and the fact that mobile clients need only support a web browser to gain access to a hotspot.

Although UAM is simple and easily deployed, it has several serious drawbacks. One problem is the user experience. Research shows that the first step to obtain network access, i.e., launching the browser, is not intuitive if the intent is to use some other application such as an e-mail client. Furthermore, enterprise users frequently require Virtual Private Network (VPN) policy settings that conflict with the requirement to access a local web server. More seriously, however, UAM typically exposes the user's credentials (username, password) to the visited network's web server—an unacceptable feature for carriers that do not wish to expose subscriber databases, even to legitimate roaming partners. Furthermore, unless the user manually inspects the certificate used by the server to secure the web pages (which is rarely done), these credentials may be unwittingly disclosed to an attacker operating a rogue wireless access point (AP).

Most of the security problems of UAM can be overcome by using Wi-Fi Protected Access, also called WPA [3]. WPA uses IEEE 802.1X [4] authentication to mutually authenticate the AP and mobile client. It also uses the Temporal Key Integrity Protocol (TKIP) to encrypt packets and prevent forgeries. The use of WPA and the seamless transition from UAM to WPA are major themes of this paper.

There are also several inconsistencies in the end-to-end architectures of currently deployed hotspots that are often caused by the variety of AAA backends of the incumbent operators. For example, broadband carriers typically use RADIUS servers natively, while Global System for Mobile Telecommunications (GSM) cellular carriers interface to SS7[2]-based backends. In a rush to offer richer, more enhanced services, service providers have also deployed a variety of proprietary systems that require

[1] http://www.wi-fi.org/OpenSection/index.asp

[2] System Signalling No. 7 is an ITU-T telecommunications standard.

complex software configurations on mobile clients. This is clearly an impediment to ease of use and ubiquitous access. Another source of inconsistency from the user's perspective is that WLAN service is sometimes provided for free. Users accustomed to free service may be confused when they discover that in a different venue they are expected to pay for service.

Tenets for a Unified Architecture

The unified hotspot architecture proposed in this paper is based on the following requirements and architectural principles:

- WLAN hotspots are essentially 802.11-based IP networks and, as such, we strongly subscribe to the use of core protocols developed in the IEEE (such as 802.1X) and the Internet Engineering Task Force (IETF). This minimizes the need for proprietary or domain-specific protocols to be used over the WLAN interface.
- The hotspot must support a common user sign-on experience that is independent of or agnostic to variations in network backends.
- The core components and interfaces of the blueprint must be agnostic to the type of hotspot operator (e.g., Wireless ISP (WISP), DSL provider, cable modem provider, or cellular operator).
- The visited hotspot must accommodate a variety of credential types (e.g., username/password, Subscriber Identity Module (SIM), and X.509 certificates) and enable new forms to be introduced over time.
- In a subscription-based access model, it must be possible to provide end-to-end security for authentication and authorization; i.e., users should be able to securely and bilaterally authenticate with the subscription home provider. The home provider should prove its identity first, before any user-specific identity is divulged, and true identities should only be exposed to the home provider.
- It must be possible for different users to avail of different levels of service depending on whether they are in the home provider's network or in a visited network.
- The framework must accommodate older UAM authentication models while articulating a clear strategy for interim coexistence and longer-term migration to more robust schemes, based on 802.1X.
- If possible, key distribution between home providers and visited networks for wireless link layer encryption should be secured and cryptographically bound to authentication and session information. (*Current* standards for WLAN key distribution do not fully meet this requirement in roaming scenarios.)
- Reauthentication when moving between access points (APs) managed by the same network operator must not cause significant delay and must not require user interaction.
- If protocol translations are required to be integrated with legacy or proprietary authentication backends, such translations should occur within the premises (architecturally speaking) of the legacy network.
- In situations where integration of services requires interworking with another network (such as a cellular operator's core data network), we advocate the notion of "loose coupling" between the WLAN hotspot and core networks. In other words, WLAN networks should be seen as standalone networks based on IEEE and IETF core protocols as opposed to radio access networks, and should not require the use of domain-specific mobility management protocols over the client's WLAN interface (for example, GPRS Mobility Management or GMM). This philosophy is also in line with a future vision where all wireless networks will be natively based on IETF's suite of IP protocols.

Generic Public WLAN Roaming Model

Figure 2 depicts a generic architecture corresponding to the hotspot blueprint for public Wireless Local Area Network (WLAN) roaming. Real-world roaming scenarios can encompass a large number of possible scenarios and network configurations. To make this complexity manageable, we define a generic roaming model that ignores non-essential aspects of roaming. For example, although a home provider may often operate a hotspot, the essential characteristic of the home provider in our model is that it maintains the user/subscriber relationship and implements an Authentication, Authorization, and Accounting (AAA) service to authenticate roaming users. Home providers do not need to provide wireless network services to fulfill this role.

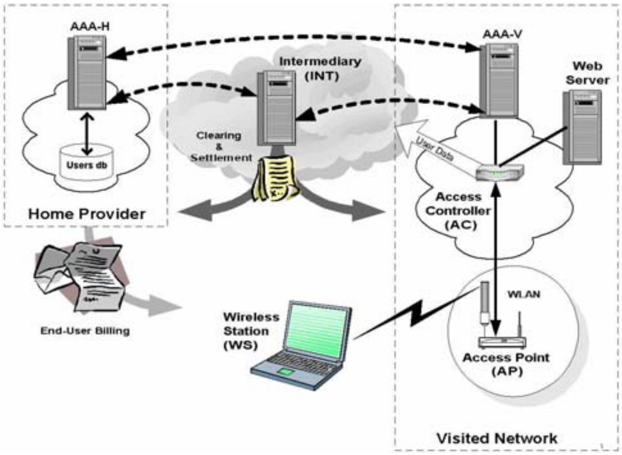

Figure 2: Generic roaming model

There are seven primary components in the generic roaming architecture:

1. The *Wireless Station (WS)* represents the user's equipment (typically a laptop computer, cell phone, or PDA) that is used to access the 802.11 network.
2. The 802.11 *Access Point (AP)* terminates the air (radio) interface to and from the WS.
3. The *Access Controller (AC)* is the entity that verifies authorization and enforces access control for authenticated users and segregates traffic of non-authenticated (guest) users.
4. The *Visited Network AAA Server (AAA-V)* serves as an AAA proxy for roaming (foreign) customers.
5. The *Home Provider AAA Server (AAA-H)* serves as the RADIUS server authenticating the WS user. The home provider and visited network operator AAA servers also participate in transactions involving the reconciliation of billing and settlement records—both online and offline—and either mutually, or via an intermediate settlement entity.
6. The *Web Server* is an optional component that could serve one or more of the following functions: browser-based login portal, local value-added services portal for guests and authenticated users, portal for new subscriptions, and redirector for other services.
7. The *Roaming Intermediary (INT)* represents a wide variety of AAA and billing intermediaries. Such functions might include AAA aggregation, wholesale hotspot service aggregation, AAA brokers and charging, billing and settlement clearing houses. They are typically implemented across multiple physical components.

It is important to note that these components are logical entities rather than literal components.

Figure 2 conceptually depicts only one possible billing model, where the home provider delivers a bill to the user. Equally valid are models where billing reconciliation is between an intermediary and the home provider or between the intermediary or visited network and the home provider.

AUTHENTICATION AND SECURITY

One of the biggest barriers to WLAN deployment is security. It is important to understand that the threats associated with network impersonation on Wireless Local Area Networks (WLAN) is substantially worse than with most other networks. With wired networks, the user's direct connection to the network has at least some level of implied authenticity by virtue of physical wires or use of virtual circuits (e.g., ATM virtual circuits) or physical circuits (e.g., dial-up). In a WLAN, there is no such first line of defense. Unless robust mechanisms to authenticate the network are employed, the user is highly vulnerable to man-in-the-middle or rogue access point (AP) attacks on the wireless link.

The sensitivity and the value of data stored on many WLAN client devices such as laptops, and the high bandwidth of WLANs offer a significant incentive to attackers. A rogue AP, since it has complete control over the channel of information flow, can perform a wide variety of attacks including eavesdropping, message insertion, message modification, Domain Name System (DNS)-based attacks, etc. Link-level encryption does not protect against this class of attacks if the attacker is one of the endpoints of the encrypted channel.

There are two basic strategies to defend against rogue AP attacks. One is to tunnel all traffic through the rogue AP using a Virtual Private Network (VPN) client and a client-hosted firewall. If executed properly, this defense limits the rogue AP to denial-of-service attacks. However, the VPN approach requires a VPN infrastructure in the network and on the client, plus robust configuration of the client firewall. These are non-trivial requirements. An alternative strategy is for the client to authenticate the network and refuse to connect to a rogue AP. Note that the latter approach is only effective if subsequent use of the connection is cryptographically bound to the authentication.

Wi-Fi Protected Access (WPA)

The security solution defined for the initial 802.11 standard called Wired Equivalent Privacy (WEP) had several documented vulnerabilities, inhibiting its use and prompting the use of UAM with VPN. The IEEE 802.11 Task Group i (TGi) has addressed these weaknesses in the new standard branded by the Wi-Fi Alliance as WPA. WPA is based on the IEEE 802.1X authentication framework, but it improves on WEP by using dynamic per-user encryption keys and per-message integrity protection. TKIP, which is used by WPA, also constructs a new per-packet encryption key in a way that defeats the Fluhrer-Mantin-Shamir attack on WEP. WPA will be eventually superceded by a TGi specification that will essentially include support for the Advanced Encryption Standard (AES) and solutions for inter-AP roaming.

With 802.1X, the WS can initially access only the unauthenticated port on the AP (or network switch behind the AP, depending on the implementation). The unauthenticated port typically limits the WS to using the Extensible Authentication Protocol (EAP) [5] protocol and communicating with the network's authentication infrastructure. If the WS and network successfully authenticate and satisfy each other's access control requirements, the session key derived by the WS is granted access to the authenticated port. The corresponding key for the WLAN network infrastructure is also communicated by the AAA-H to the AP. At this point, the WS is typically given access to the Internet.

Figure 3 depicts a typical protocol stack for WPA-based authentication. The framework permits an AP to block all unauthenticated traffic from accessing the Internet or other service networks, until the mobile client is authenticated by a provider—the visited network in prepaid or pay-for-use billing models and the home provider in subscription-based billing models.

The WPA framework relies on the EAP as the framework to carry protocol messages between the supplicant (client), authenticator (AP), and authentication server. The EAP messages are carried over EAPOL (EAP over LAN) frames between the WS and the AP and reencapsulated in RADIUS messages from the AP to the home AAA server (via zero or more AAA proxies). For security reasons, RADIUS is sometimes also carried over IPsec. In future, RADIUS may be natively substituted by DIAMETER, a successor AAA protocol being developed by the IETF.

The WPA/802.1X model offers significant advantages over the browser hijack model. One of the most important advantages is that 802.1X is designed to support extensible end-to-end authentication between the WS and the home provider's AAA-H.

The EAP channel established by 802.1X can support a variety of authentication protocols and credential types—all that is needed is an EAP method with appropriate security properties describing how the protocol is encapsulated by EAP. The EAP method must (a) perform mutual authentication, (b) derive fresh session keys, (c) be immune to man-in-the-middle attacks, and (d) be immune to dictionary attacks.

When the EAP channel is established between the WS and the AAA-H, there is no need for the visited network's AP, AC, or AAA-V to comprehend or support the specific EAP method or credential types used by the home provider. This feature provides great flexibility to the client and service providers.

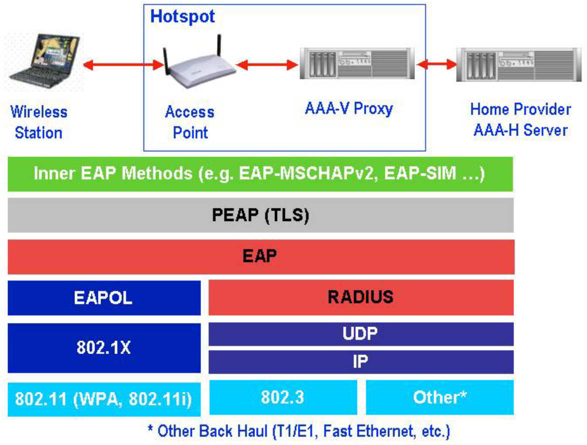
Hotspot
Wireless Station
Access Point
AAA-V Proxy
Home Provider AAA-H Server
Inner EAP Methods (e.g. EAP-MSCHAPv2, EAP-SIM ...)
PEAP (TLS)
EAP
EAPOL
RADIUS
802.1X
UDP
IP
802.11 (WPA, 802.11i)
802.3
Other*
* Other Back Haul (T1/E1, Fast Ethernet, etc.)

Figure 3: 802.1X/WPA with PEAP

To achieve end-to-end identity confidentiality, it is recommended that the Protected EAP [6] tunneling authentication protocol be used. PEAP is an authentication tunneling protocol that creates a protected channel for other EAP-based authentication methods. PEAP enables two-phase mutual authentication where the network first authenticates to the client via a digital certificate and then the client authenticates to the network using some other EAP method inside an encrypted channel. The client authentication method is typically based on passwords rather than client certificates. This is the same model used by secure web sites; however, the use of domain-specific root certificates with PEAP greatly improves the trust model over the more traditional browser used by e-commerce. This is because the large number of commercial root certificate authorities trusted by browsers has a business incentive to issue as many certificates as possible, and it is relatively easy for an attacker to obtain such a certificate.

With PEAP, common session key derivation, distribution, and configuration solutions can be defined for a variety of credential types, including certificates, username and password, and Universal Subscriber Identity Modules USIM. If industry alignment can be achieved in these areas, it will be easier for network operators to support a variety of roaming scenarios across different network types. PEAP helps address the user security requirements in the 3GPP Release 6 TR 22.934 document. Furthermore, the TLS and the PKI infrastructures used by PEAP can also help address many of the requirements for network operator security features listed in TR 22.934. Alternatives such as Tunneled TLS [7], which has similar functionality, can also be used without significantly sacrificing interoperability, because of the end-to-end properties of EAP. However, PEAP is likely to be more widely deployed on client platforms due to native operating system integration.

A new version of PEAP is being developed that includes a fix for man-in-the-middle attacks that are possible when the same credentials are used both inside and outside of PEAP. When this version of PEAP becomes available, we recommend only using it with inner EAP methods that can derive keys. In the case of passwords, this would imply the use of EAP-MSCHAPv2 over legacy methods such as EAP-OTP, EAP-MD5, or EAP-GTC. Also, while PEAP may seem redundant when used with EAP methods that inherently offer bilateral authentication and 128+ bit key derivation (like EAP-AKA with USIM cards), PEAP has a property called session reuse that can optimize handoffs across APs.

IP and MAC Address Filtering

A common method for controlling internet access for WLAN networks is to filter packets based on the source IP address and/or MAC address. This method can be used to limit a WS access to only designated destination addresses such as the browser hijack web server. Although this is a very common method for access control, in many cases proprietary implementation methods are used. This is an area where additional standardization work may be needed.

VLAN Support

Virtual LAN (VLAN) technology can provide additional flexibility to the deployment of WLANs, because it can be used to provide logical isolation of traffic sent through common WLAN APs. This can be used to enhance the security of portions of the traffic to support more robust billing methods, enable infrastructure sharing among carriers, and to separate private WLAN connections from public access traffic. For example, broadcasts directed to secure network segments are encrypted and thus protected from weakly authenticated users, if the traffic is separated. In 802.11, there are no physical VLAN ports, so VLAN membership is often assigned dynamically as part of the authentication process via RADIUS accept messages. Another possibility is to assign VLANs one-to-one with 802.11 Service Set Identifiers (SSIDs). VLANs can also provide a mechanism for associating users with site-to-site tunnels used to direct data traffic to the core networks of roaming partners. VLANs can be used in conjunction with IP and MAC address filtering to control what parts of the network are available to specific WSs.

Migration to WPA

Although we prefer 802.1X and WPA, we also recognize that until 802.1X-capable clients are widely deployed, there will be a market requirement to support the Universal Access Method (UAM). Furthermore, even when 802.1X is used, browser hijacking can be useful to help resolve authentication failures and to permit the establishment of new accounts. Therefore, the generic hotspot architecture supports a mixture of UAM and 802.1X-based authentication.

Figure 4 illustrates a possible coexistence strategy involving the use of a VLAN-capable AP to separate UAM traffic from 802.1X traffic. To support both 802.1X and UAM, each AP supports two different Service Set Identifiers (SSIDs), one corresponding to 802.1X and one open (for UAM). With current AP hardware, only one of these SSIDs would be advertised by the AP (corresponding to WPA), but the other (for UAM) could be discovered via the 802.11 probe request/response mechanism. The open SSID would not require any link-layer security, but the AC would limit user access to the local web server until the user obtains authorization to use the network. Subsequent enforcement of access control for the UAM method is likely to be based on the client's MAC address, which is not very robust. Attackers can easily configure their own equipment with the same MAC address and masquerade as legitimate users, stealing their bandwidth. This creates a business incentive for network providers to migrate users away from the UAM as soon as possible.

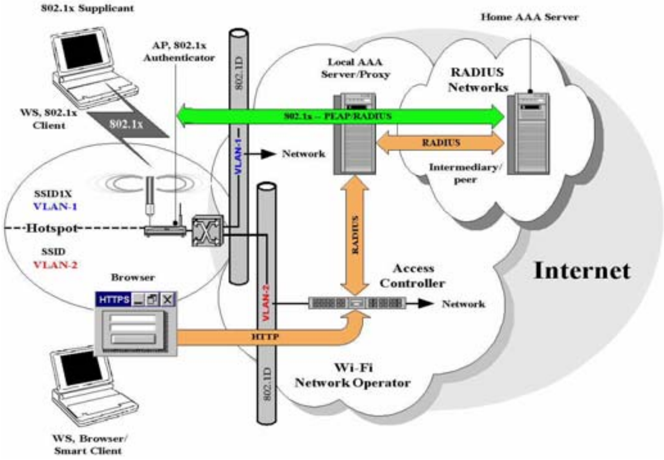

Figure 4: 802.1X and UAM/browser hijack coexistence

If an AP is capable of advertising multiple SSIDs, the WS will be able to detect them and choose the appropriate one. If the AP can only broadcast one SSID, the WS may be able to probe for the hidden one by using a database of hidden SSIDs. Since such preconfiguration is not practical when roaming, the industry needs to develop better ways to discover hidden 802.11 SSIDs.

The AP assigns separate VLAN tags to packets according to the SSID the WS is associated with. The VLAN switch in turn routes the packets during authentication so that the 1X traffic gets sent to the AAA-V for authentication, and the browser hijack mechanism is used for the non-1X clients. The web server for the browser hijack is not shown in the figure (it could be implemented by the access controller). Note also that although the figure does not show the access controller in the data path for the 1X traffic, in many implementations it will be.

Other coexistence models are possible as well. For example, if traffic from the 802.1X clients and browser hijack clients is mixed, the VLAN switch can be eliminated, and the Access Controller can manage both types of clients. However, this approach is not as secure as the approach shown in Figure 4.

ARCHITECTURAL CONSIDERATIONS FOR ONE-BILL ROAMING

Users of hotspot services could participate in one of several billing models; prepaid, pay-for-use, and postpaid (subscription-based) are likely to be the most common. Furthermore, charging itself could be based on fixed or flat rates, based on usage (time, bytes and/or number of connections) and services used. Regardless of the billing model, roaming users should have the same experience when connecting to a visited network as they do when connecting to their home network. Ideally, charges associated with WLAN roaming usage would appear in an integrated single bill as is the case for cellular voice roaming today.

Prepaid and pay-for-use settlement procedures are often localized to the visited operator or managed by a clearing house on behalf of the visited operator. Postpaid billing, on the other hand, requires business agreements between the visited and home operators. The simplest scenario is one in which each operator executes a bilateral agreement with every peer roaming partner. In the world of public WLANs, this may not be appropriate for two reasons:

- The number of service providers will be very large, creating a scalability problem.
- Since incumbent operator communities subscribe to different billing and settlement practices, there may be incompatibilities.

This is where clearing houses can play a role. Figure 5 conceptually depicts a Data Clearing House (DCH) serving the role of a settlement intermediary. The DCH connects via a chain of RADIUS/AAA proxy intermediaries to one or more hotspots and transacts charging, rating, billing, and settlement with diverse backends, using protocols such as RADIUS, TAP, CIBER, IPDR, and others. DCHs specialize in being able to convert different record formats to one format and in providing other value-added services such as re-rating and fraud detection.

For example, different versions of TAP may be supported by the billing subsystems of different home providers. A DCH can do the necessary conversions so that accounting or charging records originating from a visited WLAN can be processed correctly by a given home provider.

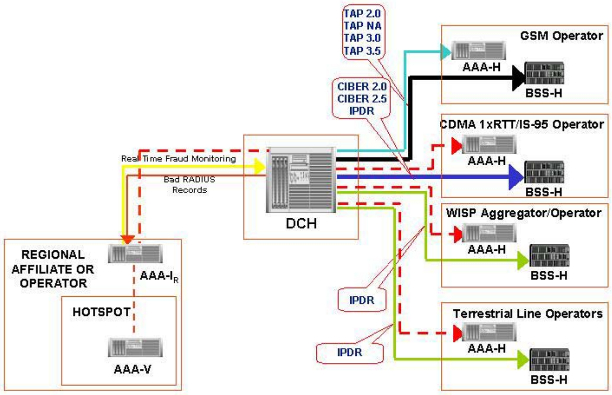

Figure 5: WLAN Inter-operator postpaid settlement

For one-bill roaming to work on a global scale, we propose a two-tiered model as depicted in Figure 6 below. This notion of a backbone of roaming intermediaries (only one DCH is shown in the figure for simplification) results in better scaling of roaming agreements and RADIUS traffic aggregation. Such a system can also combine prepaid and postpaid billing models.

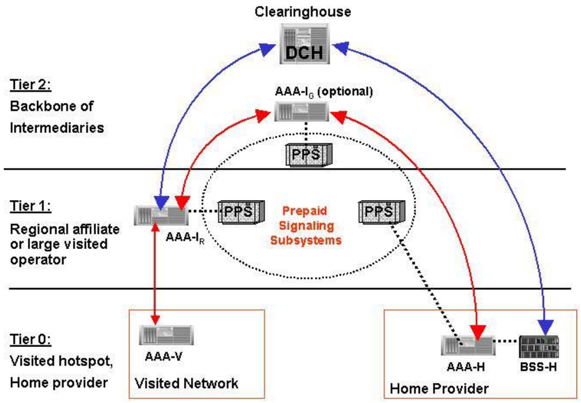

Figure 6: Two-tier model for billing and settlement

The idea of using regional affiliates in conjunction with global clearing houses has proven to be an effective model for billing and settlement scalability in the cellular industry.

THE ROAD AHEAD

Future research will focus on the following areas as this architecture continues to evolve to support more advanced usage models:

- *Fast and seamless inter-access point (AP) handoffs*: The next big step is the introduction of value-added services to public hotspots that will include messaging, real-time multimedia streaming, and data application portals. These services require fast and lossless handovers across APs. The suite of solutions include improvements in TGi for pre-authentication and fast reauthentication, early Protected EAP (PEAP) termination and PEAP session resumption, as well as secure context transfers across APs

- *Wireless wide area network (WWAN) interworking*: Authorization and access to IP data services in GSM and CDMA core networks (2.5G and beyond) from WLAN hotspots will require solutions such as mobile IP overlay over GPRS for IP session persistence across WWAN and WLAN, the use of tunneling to ensure IP address reachability, and enabling of access to native IPv6 services like IP multimedia subsystems (IMS) over heterogeneous IPv4-IPv6 clouds. There are also significant challenges inherent in services provisioning and authorization, given the diverse set of operator types.

- *Public key-based authentication and authorization*: The use of public key-based authentication with attributes for dynamic services provisioning and authorization will overcome cryptographic limitations with use of passwords, not require use of expensive legacy token schemes like Generic Token Card (GTC) and SIM, and promote a more homogeneous framework for network access, whether in the home, enterprise, or public hotspots.

- *Network and services discovery*: There is a need to create a common yet extensible standardized framework for hotspot discovery, selection of service providers, and provisioning and use of services.

SUMMARY

In this paper, we examined a variety of issues related to public Wireless Local Area Network (WLAN) roaming. Although there is substantial interest in developing a global WLAN roaming market, unless the technical and business challenges are addressed in a coordinated manner, what is more likely to emerge is a fragmented and incompatible tangle of proprietary solutions and regional alliances with no easy path to convergence.

Although many issues are still unresolved, the industry should at least rally behind the following strategically important points:

- Standards-based solutions should be used whenever practical.

- Authentication should migrate to 802.1X and WPA, which have superior security properties and permit greater flexibility in credential types than browser hijack. However, both browser hijack and 802.1X will coexist for at least a few years.

- Accounting data suitable for all billing models should be collected.

- Roaming intermediaries such as aggregators and clearing houses will help solve scalability issues and provide interoperability with legacy authentication and billing systems.

The hotspot blueprint architecture derived from this study will be implemented and tested in a validation test bed by a variety of carriers and vendors in the Asia Pacific region. Results from the test bed, including feedback from participating carriers, vendors, and other interested parties, will be used to develop specific deployment recommendations for WLAN client vendors, hotspot operators, and AAA providers.

ACKNOWLEDGMENTS

The authors thank Jesse Walker, Lakshmi Ramachandran, and Sameer Pareek for their contributions to this paper.

GLOSSARY OF ACRONYMS

This glossary will help the reader navigate the many acronyms used in this paper.

Acronym	Definition
AAA	Authentication, Authorization, and Accounting
AAA-H	Home provider AAA server
AAA-V	Visited network AAA server/proxy
AP	access point
ATM	Asynchronous Transfer Mode
AES	Advanced Encryption Standard
CIBER	Cellular Intercarrier Billing and Exchange Roaming Record
DCH	Data Clearing House
DNS	Domain Name Service
EAP	Extensible Authentication Protocol
GSM	Global System for Mobile Telecommunications
ISP	Internet Service Provider
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IMS	IP Multimedia Subsystem
INT	Roaming intermediary
IPDR	Internet Protocol Detail Record
LAN	Local Area Network
PEAP	Protected EAP
RADIUS	Remote Access Dial-In User Service
SIM	Subscriber Identity Module
SSID	Service Set Identifier, a unique 32-bit identifier for WLANs
TAP	Transferred Account Procedure
TKIP	Temporal Key Integrity Protocol
TLS	Transport Layer Security
TTLS	Tunneled TLS
UAM	Universal Access Method
USIM	Universal Subscriber Identity Module
VLAN	Virtual LAN
VPN	Virtual Private Network
WEP	Wired Equivalency Privacy
WISP	Wireless ISP
WLAN	Wireless Local Area Network
WPA	Wi-Fi Protected Access
WS	Wireless Station

REFERENCES

- [1] "Enterprises Must Consider Six Types of WLAN 'Hot Spots'," Gartner, June 2003.
- [2] Transport Layer Security, published as IETF RFCs 2716 and 3546.
- [3] Wi-Fi Protected Access, http://www.wi-fi.org/OpenSection/protected_access.asp
- [4] The IEEE 802.1X standard, http://standards.ieee.org/getieee802/download/802.1X-2001.pdf
- [5] Extensible Authentication Protocol, IETF RFC 2284bis (ongoing).
- [6] Protected EAP, IETF draft-josefsson-pppext-eap-tls-eap-06.txt (ongoing).
- [7] Tunneled TLS, IETF draft-ietf-pppext-eap-ttls-02.txt.

AUTHORS' BIOGRAPHIES

Prakash Iyer is a senior staff architect in Intel's Network Architecture Lab. In his 11+ years at Intel, his areas of focus have included 10/100 Mbps Ethernet, IP Telephony and video conferencing, IPv6, networking security, and residential networking. He currently leads a research team focused on advanced wireless technologies. Prakash was the chair of the UPnP Internet Gateway group in the UPnP Forum and was a 2002 recipient of the Intel Achievement Award for his UPnP work. He holds B.S. degrees in Physics and Electrical and Computer Engineering and an M.S. degree in Computer Science. His e-mail is prakash.iyer@intel.com

Victor Lortz is a software architect in Intel's Network Architecture Lab. In his nine years at Intel, his areas of focus have been object-oriented software development, home networking, network security, and wireless mobile networking. Victor is the chair of the UPnP Security Working Committee. He was a 2002 recipient of the Intel Achievement Award for his UPnP work. He holds a B.A. degree in Physics from Whitman College and M.S. and Ph.D. degrees in Computer Science from the University of Michigan. His e-mail is victor.lortz@intel.com

Lee Tapper is a program manager and business development manager in Intel's Emerging Platform Lab. Lee works on 802.11 programs and was a member of the GSMA's WLAN task force. Lee holds a BSEE degree from the University of Texas. His e-mail is lee.s.tapper@intel.com

Roger Chandler is a market development manager in Intel's Emerging Platform Lab. His areas of focus have included manufacturing process analysis, high-performance 3D technologies, and home networking. He was a 2001 recipient of an Intel Achievement Award for his work in the field of Web 3D. He is currently focused on market development strategies for Intel's many wireless technology initiatives. He holds an MBA degree from the University of Georgia and a B.A. degree from the University of Tennessee. His e-mail is roger.d.chandler@intel.com

Roxanne Gryder is a business development manager in the Networking Architecture Lab, Intel R&D. She is involved in business and marketing activities for a variety of wireless technologies in the labs. She received B.S. and MBA degrees from Northwestern University and Cornell University, respectively. Her e-mail is Roxanne.r.gryder@intel.com

On the Union of WPAN and WLAN in Mobile Computers and Hand-Held Devices

Ofer Bar-Shalom, Wireless Communications & Computing Group, Intel Corporation
Gordon Chinn, Mobile Platforms Group, Intel Corporation
Kris Fleming, Mobile Platforms Group, Intel Corporation
Uma Gadamsetty, Mobile Platforms Group, Intel Corporation

Index words: Wireless, Bluetooth, WPAN, WLAN, 802.11, Centrino^{TM}, Intel PCA, Coexistence

ABSTRACT

With the introduction of Intel® Centrino^{TM} mobile technology and Personal Internet Client Architecture (PCA), today's mobile notebooks and mobile devices are advanced communication platforms with multiple wired and wireless technologies. Of those wireless technologies, Wireless Personal Area Networks (WPANs) are a key ingredient for current and future versions of Intel Centrino mobile technology and PCA. This paper focuses on how WPAN technology such as Ultra Wideband (UWB) or Bluetooth* technology enables new usage models when used in conjunction with Wireless Local Area Networks (WLANs) technology. Several key usage models of WPAN and WLAN interaction are explored in detail. In addition to the proposed simultaneous use of WPANs and WLANs, we elaborate on integration and wireless coexistence issues for WPANs and WLANs. Both radios operate in the 2.4 GHz Industrial Scientific Medical (ISM) band and therefore interfere with each other. We explain how Intel's coexistence solution diminishes the potential interference issues that result from this simultaneous use of multiple wireless technologies.

We explore how the new usage models are enabled by the union of WLAN and WPAN. We scan the WLAN and WPAN environment and highlight some of the major new advances. For WPAN, we review features in the upcoming 1.2 version of the Bluetooth specification as well as the formation of the IEEE 802.15.3a Working Group (WG). For WLAN, we explore the rapid advances in security and Quality of Service (QoS) definitions of IEEE 802.11 WGs, and the new IP protocols that support mobility as defined by the Internet Engineering Task Force (IETF). These new usage models drive improved wireless capabilities for mobile platforms and introduce new scenarios that the mobile user can take advantage of with Intel's next-generation notebooks and hand-held devices.

INTRODUCTION

Intel Centrino mobile technology and Personal Internet Client Architecture (PCA) are key ingredients for enabling the union of Wireless Local Area Networks (WLANs) and Wireless Personal Area Networks (WPANs) and the new usage models enabled by this union. In the first part of this paper we briefly review the technology behind Intel Centrino mobile technology and PCA and we give a short primer on the main wireless technologies that are discussed in the paper. In the second section, we examine how the wireless networking capabilities of Intel Centrino mobile technology and PCA enable new usage models. We then explore the various technology drivers that are influencing these new usage models. In the final section we address platform integration issues and coexistence solutions.

Intel Centrino Mobile Technology Overview

Intel Centrino mobile technology combines the Intel® Pentium® M processor, the Intel® 855 chipset family, and

®Intel and Centrino are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

®Intel and Pentium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

the Intel® PRO/Wireless 2100 Network Connection, as shown in Figure 1.

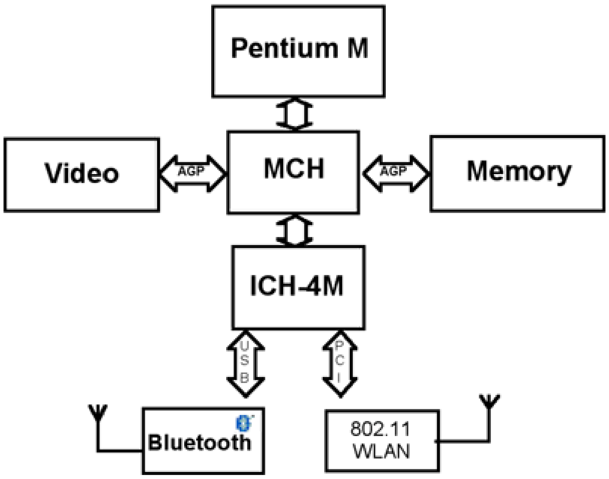

Figure 1: Intel® Centrino™ mobile technology

The Intel Pentium M processor micro-architecture is optimized for high performance and low power. The Intel 855 chipset family consists of two components: the Intel® 855PM Memory controller Hub (MCH) and the Intel® 82801DBM I/O Controller Hub (ICH-4M). The Intel PRO/Wireless 2100 Network Connection is the integrated Wireless LAN (WLAN). All components were designed, optimized, validated, and tested to work together with mobility in mind. Many Intel Centrino mobile technology laptops are adding WPAN capabilities. With the addition of WPAN, new usage models are enabled and integration issues need to be considered. These new usage models and integration issues are addressed in this paper.

The Intel Personal Internet Client Architecture

The Intel PCA [15], [16] partitions the device configuration of the traditional cellular platform into an Applications Subsystem (APPS), a Communication Subsystem (COMM), and a Memory Subsystem. This partitioning allows application development to evolve independent from communication standards. The Intel PCA provides open programming interfaces and services between physical platforms (including communication) and application software, thereby facilitating faster development of the application and abstracting the underlying physical resources.

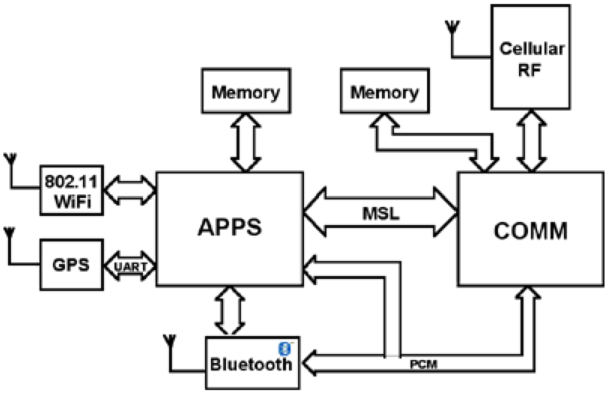

Figure 2: PCA with wireless subsystems

The COMM is based on the Intel® Micro Signal Architecture digital signal processor (DSP), and it provides the APPS the services to access cellular wireless networks that are independent of the physical medium. The COMM services are independent of the air-link technology, and they are responsible for maintenance of the connection to an appropriate wireless network for telephony and data services in support of the APPS.

The APPS is based on Intel® XScale™ technology and is capable of running an operating system, user interface, and applications. It manages resources such as user input/output devices, expansion devices, memory interfaces, power management, and communication interaction with the COMM. In addition, the APPS processor hosts the wireless subsystems including the Bluetooth, 802.11, and the Global Positioning System (GPS), which provide various kinds of data connectivity. All of these wireless technologies are described in the next section.

Figure 2 describes a block diagram of an implementation example of a cellular mobile system in PCA architecture. The COMM has an RF connection to the cellular system and is connected to the APPS using a standard link called the Intel® Mobile Scalable Link (MSL). The MSL hardware is capable of running up to 52 MHz or has a data throughput of 208Mbps, and it is used mainly for exchanging data packets between subsystems. The Bluetooth baseband+radio chip is connected using a Universal Asynchronous Receiver Transmitter (UART) connection; it also has serial Pulse Coded Modulation (PCM) interfaces to both APPS and COMM processors. The PCM connection to the APPS is used during audio scenarios handled by the APPS processor such as voice memo pad (VMP), voice recognition (VR),

® Intel and XScale are registered trademarks and trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

calendar/incoming e-mail alerts, and so on. The direct connection to the COMM processor is used during cellular voice calls. Finally, the 802.11 MAC and PHY are connected using some proprietary interface.

WIRELESS TECHNOLOGIES PRIMER

The increasing demand for a variety of high-rate data services along with the requirement for reliable connectivity anywhere at anytime demands that several different wireless subsystems be integrated into a single hand-held device or into a mobile notebook computer.

The market share of devices enabled with Bluetooth wireless technology has been growing constantly. Analysts predict that by 2005 the number of mobile handsets enabled with Bluetooth wireless technology will be as high as 800 million units with a market penetration of 65% [9]. Side by side with the Bluetooth technology, the 802.11 WLAN is also becoming the *de facto* wireless standard for high-rate Local Area Networks (LANs), with more and more wireless networks being deployed around the world every day. In addition, as more and more location-based data-services and emergency services (E-911) are becoming commercially available, the demand for GPS receiver integration in hand-held cellular devices increases accordingly.

Following, we briefly describe the main wireless technologies mentioned above, which are being considered for both PCA and Intel Centrino technology. Our intention is not to cover all the issues of a specific technology, but only highlight the main features describing how that specific technology can serve as a data gateway.

- ***Wireless cellular technologies****:* The wireless cellular systems being considered in this paper belong to the second generation (2G), 2.5G, and 3G. The majority of the cellular networks being deployed around the world is based on the Global System Mobile (GSM) standard. This is a TDMA standard that can operate in several radio frequency (RF) bands. The most usable of these bands are the 900MHz band, the 1800MHz band (a.k.a. digital communications system (DCS) band), and the 1900MHz band (a.k.a. personal communications system (PCS) band). Each of these operation bands is divided into channels in widths of 200kHz. The 200kHz channel is divided into eight time slots and is shared in a time division manner between eight different users. The GSM network is used mainly for conventional telephony services, but can also accommodate internet browsing via circuit switched dial-up networking (DUN), providing raw data rates of 9.6kbps.

The evolution for the 2.5G of the TDMA standards includes the General Packet Radio Service (GPRS) and the Enhanced Data Rates for GSM Evolution (EDGE). The GPRS standard provides a packet network on dedicated GSM channels. The GPRS retains the original modulation formats specified in the original 2G GSM standard, but uses a completely redefined air interface in order to better handle packet data access. GPRS subscriber units are automatically instructed to tune to dedicated GPRS radio channels and particular time slots for “always on” access to the network. The GPRS standard allows a single subscriber unit to occupy multiple consecutive time slots in order to increase its data throughput. When all eight time slots of a GSM radio channel are dedicated to GPRS, an individual user is able to achieve as much as 171.2kbps [20].

EDGE (which is sometimes referred to as Enhanced GPRS, or EGPRS) includes a new digital modulation format in addition to the standard GSM modulation. EDGE allows nine different air-interface formats known as multiple modulation and coding schemes (MSC), with varying degrees of error control protection, to be selected according to the instantaneous network and operating conditions. Similar to GPRS, the EGPRS standard also allows a single subscriber unit to occupy multiple time slots for increasing its data capacity. In practical network conditions, when all eight GSM time slots are being dedicated to a single user, the EDGE is capable of providing a raw peak throughput data rate of 384kbps.

Third-generation (3G) systems promise unparalleled wireless access such as Voice over Internet Protocol (VoIP) and unparalleled network capacity. The eventual third evolution for the 2G GSM networks is based on wideband code division multiple access (W-CDMA), also known as Universal Mobile Telecommunications Service (UMTS). W-CDMA assures backward compatibility with the 2G GSM networks, as well as with the 2.5G TDMA networks such as EDGE. The W-CDMA air interface was designed for “always-on” packet-based wireless services and can support packet data rates of up to 2.048Mbps.

Mobile handsets based on Intel PCA technology will support all 2G, 2.5G, and 3G cellular technologies mentioned above, and will provide constant internet connectivity with a high quality of service (QoS) throughout the entire coverage area of the cellular network, including regions that are not covered by other internet access means such as 802.11, described next.

- ***IEEE 802.11*** (a.k.a. Wi-Fi, or Wireless Local Area Network (WLAN)): This was defined by the IEEE in 1997 as a standard that would replace the wired Ethernet connection cables with a wireless connection. The 802.11 standard is limited in scope to the Physical (PHY) layer and the Medium Access Control (MAC) sub-layer, with MAC origins to the IEEE 802.3 Ethernet standard. The 802.11 standard definition for the Physical layers standard includes definitions for Physical Layer Convergence Procedure (PLCP) and Physical Medium Dependent (PMD) sub-layers. The 802.11 has three major extensions that are being considered today. The first extension is called 802.11a. It operates in the 5GHz band with twelve separate non-overlapping channels. As a result, you can have up to twelve access points (APs) set to different channels in the same area without them interfering with each other. The 802.11a uses orthogonal frequency division multiplexing (OFDM) PHY, which divides a data signal across 48 separate sub-carriers within a 20MHz channel to provide transmission rates of 6, 9, 12, 18, 24, 36, 48, or 54Mbps.

  The second extension is called 802.11b and is the basis of the majority of wireless LANs in existence today. The 802.11b operates in the 2.4GHz ISM band and uses direct sequence spread spectrum (DSSS) PHY with complementary code keying (CCK) modulation to disperse the data frame signal over a relatively wide (approximately 30MHz) portion of the 2.4GHz frequency band. The data rates supported by the 802.11b are 1, 2, 5.5, and 11Mbps.

  The third extension, which has recently been approved, is called 802.11g. 802.11g broadens the 802.11b's data rates to 54 Mbps within the 2.4 GHz band using OFDM, similar to the one being used by the 802.11a. Similar to 802.11b, 802.11g operates in the 2.4GHz band, and the transmitted signal uses approximately 30MHz, which is one third of the band. This limits the number of non-overlapping 802.11g APs to three, which is the same as 802.11b.

- ***GPS:*** A growing number of hand-held devices integrate Global Positioning System (GPS) receivers in their design. GPS is used for both security applications such as E-911 and location-based services. Similar to the 802.11b, the GPS also uses the DSSS RF technology that operates in the spectrum band of 1570-1580MHz. The GPS receiver measures the time of arrival (TOA) for signals transmitted from 24 satellites, which are in orbit around earth. The TOA data from three or more satellites is used for calculating the instantaneous location of the GPS receiver. The acquisition of the satellite signals usually requires a line of sight (LOS) between the antenna and the transmitting satellite thereby making GPS less applicable for indoor locations. The signal acquisition is sometimes quite lengthy and when it is, it reduces the idle periods of the system incorporating a GPS receiver. In order to shorten it and increase the robustness of the system for indoor operations, it is often used to integrate a variant of the GPS called "assisted GPS." With assisted GPS, the mobile phone receives initial information about the satellites above it and uses this information to reduce the acquisition time and improve its GPS reception sensitivity.

- ***Bluetooth Technology:*** The Bluetooth specification defines a frequency-hopping spread spectrum (FHSS) wireless system, operating in the 2.4GHz band. It was designed as a wireless cable replacement and as a Personal Area Networking (PAN) technology. The current revision of the Bluetooth specification (1.1) supports data rates as high as 723kbps and can accommodate (unlike 802.11 technology) both asynchronous (i.e., packet switched) data and synchronous (i.e., circuit switched) audio transmissions. More information on Bluetooth technology is found in the coming sections.

INTEL CENTRINO MOBILE TECHNOLOGY AND PCA USAGE MODELS

With Intel Centrino mobile technology and the Intel Personal Internet Client Architecture (PCA), today's mobile notebooks and mobile devices are advanced communication platforms and can interact with one another to enable common usage models. Additionally, these communication advancements are due to the integration of Wireless Local Area Network (WLAN)- and Wireless Personal Area Network (WPAN)-enabled mobile notebooks with Intel Centrino mobile technology and are also due to mobile phones/PDAs that use PCA so that new usage models are relevant to each platform.

Common Usage Models

One of these key common usage models is the interaction of the mobile notebook with the mobile phone. In this model, the mobile phone provides access from its Wireless Wide Area Network (WWAN) to the mobile notebook. The mobile notebook and mobile phone use their WPAN radio to wirelessly exchange data. The mobile phone acts as an access point (AP) (or a modem) to the mobile notebook to share its General Packet Radio Service (GPRS), Enhanced GPRS (EGPRS), or 3G data connection with other mobile devices to access the Internet.

Another key usage model of mobile notebooks and mobile phones is the use of the mobile notebook as a data store. In this model the mobile phone is used to access information stored on the laptop. For example, calendar, contacts, and mail information stored on the mobile notebook can be accessed and/or can synchronize the mobile phone via the WPAN interface. Additional information such as maps, media files, and pictures can be accessed by using the same mechanisms.

Intel Centrino Mobile Technology Usage Models

With these communication advancements in Intel Centrino mobile technology, several new usage models are more pertinent for the mobile notebook. One example is the use of the mobile notebook as a "mobile router." A mobile router is a device with multiple networking interfaces that routes network packets via multiple network interfaces just as a normal router would, but the difference is that the mobile router moves from network to network. Notebooks with Centrino Mobile Technology with WLAN and WPAN capability are ideal mobile routers. The mobile notebook uses the WPAN interface to connect a user's personal devices (phone, PDA pagers, watches, etc.) to the Internet via other networking interfaces. The mobile notebook uses its WLAN connection to connect to a wireless hotspot to provide internet access to the user's personal devices. As the mobile notebook moves from hotspot to hotspot, internet connectivity is maintained for the user. Personal devices, which are connected by a WPAN connection, are not directly affected by the movement because they maintain the same WPAN connection with the mobile router and move with the mobile router. The mobile router performs the necessary tasks to maintain connectivity as the mobile router moves from hotspot to hotspot.

Another key usage model is the use of Voice over IP (VoIP) with Centrino mobile technology. The integration of Intel Centrino mobile technology with VoIP will allow the mobile notebook user to make and receive voice calls. The WLAN interface is used to provide connectivity to the Internet to send and receive the VoIP packet. The WPAN interface is mainly used for cable-replacement applications such as connecting to an audio headset while the phone acts as an audio headset to send and receive the audio for the VoIP session.

PCA: Key Wireless Usage Models

Data is the name of the game today and in the future. Any design that will be able to provide high-rate, versatile, and reliable connectivity will be the design of choice. The essence of the PCA is to allow data streaming by various means, simultaneously and without mutual interference between the involved subsystems. Data from the outside world can stream into the system through either the cellular link (i.e., Global System Mobile (GSM), GPRS, EGPRS, and the Universal Mobile Telecommunications Service (UMTS); through 802.11, Bluetooth link (e.g., PAN, downloaded JAVA applications enabled with Bluetooth technology, etc.), Infra Red Data Association (IrDA), serial-cable connections (e.g., the Universal Asynchronous Receiver Transmitter (UART), the Universal Serial Bus (USB), Secure Digital (SD) cards, etc.), and a Global Positioning System (GPS). Any data, whether they are downloaded through the cellular link, or via any peripheral subsystems connected directly to the Applications Subsystem (APPS) processor will be processed by the APPS processor. The reason for the partitioning is twofold. First, it ensures that downloaded applications (e.g., JAVA games) can be executed without overloading the Communication Subsystem (COMM) processor. Second, it disables any ability of the application to harm the execution of the COMM processor, (e.g., by Trojan horse virus attacks, by memory leakages as a result of poor application design, etc.). Overall, this partitioning enables continuous connectivity to the cellular network to be maintained.

The variety of connectivity means, together with the requirement for simultaneous coexistence of those means, enables the support of hybrid usage models combining data and voice. In the following section, we review some usage scenarios and briefly describe how the wireless subsystems are being used in each of them. Speaker independent voice recognition (SIVR) can be used for hands-free dialing, smart phonebook browsing, and so forth. Natural accessories for this kind of application are hands-free devices, enabled with Bluetooth technology, such as headsets and wireless car kits, for which the PCA device serves as an audio gateway (AG). As an example, a scenario that combines voice with data is video conferencing. In this scenario, video data from the remote site can stream into the system either through the cellular network or via an 802.11 hotspot, located in local coffee shops, at the office, or at an airport. The video data can be either processed by the APPS processor for display on the PDA screen. Alternatively, it can be routed into the Bluetooth protocol stack for streaming into a projector enabled with Bluetooth technology, thereby enabling it to be displayed on a bigger screen. The video and audio of the near end can stream from a digital camcorder, enabled with Bluetooth technology, into the APPS processor and from there either over cellular or WLAN to the other side. The audio can be processed by the APPS and sent over the Bluetooth link right into the user's Bluetooth headset. Another usage example of Bluetooth technology that is becoming more and more popular involves location-based information. A user's location can be measured either by GPS, WLAN hotspots, or even by roadside Bluetooth kiosks, and information can be provided on such things as traffic jams, retail sales in malls on the user's route, etc. Clearly, this information would have to be retrieved while the user is on a voice call using his or her Bluetooth headset or car-kit, so some coexistence mechanism that enables the cellular, the Bluetooth technology, the WLAN, and the GPS is required. Such a mechanism is described later in this paper.

TECHNOLOGY-ENABLING DRIVERS

During the past couple of years several key technologies have been developed or are in the process of being developed that help to enable the new usage models described in previous sections. These key technologies are major driving factors in enabling the union of Wireless Local Area Networks (WLANs) and Wireless Personal Area Networks (WPANs). Of these new technologies, we review four in this paper: the IEEE 802.11i standard, the 1.2 version of the Bluetooth SIG specification, Ultra Wideband (UWB) (IEEE 802.15.a) technology, and the Internet Engineering Task Force (IEFT) mobility drafts. Each of these key technologies is a major driver in an effective union of WLANs and WPANs.

IEEE 802.11i

The IEEE 802.11 security task group (Tgi) is developing specification improvements to increase the security of the current 802.11. There has been a high market demand for an interoperable WLAN security solution, and Wi-Fi Protected Access (WPA) is the wireless industry's response. WPA is a subset of the work from the IEEE 802.11 security task group and addresses the vulnerabilities that the Wired Equivalent Privacy (WEP) defined in the 1999 standard; it does not require a hardware upgrade to support it. WPA is a transitory step between the original WEP security solution and the forthcoming IEEE 802.11i standard. WPA provides additional security mechanisms to ensure that enterprise and consumer wireless users' data are protected. This is the first crucial step in enabling the new usage models described in the previous usage models.

The Bluetooth 1.2 Specification and New Profiles

The 1.2 version of the Bluetooth SIG specification, scheduled to be formally adopted by the end of Q3 '03, includes some new features as well as enhancements of existing features from the 1.1 specification. The improvements over the 1.2 specification are faster connection establishment, adaptive frequency hopping (AFH), extended Synchronous Connection Oriented (SCO) links (eSCO), scatter mode (for enhancing the PAN working mode), anonymity mode (for improved user's privacy, in coordination with the emerging location technologies enabled with Bluetooth technology), and some improvements in the Logical Link Control and Adaptation Protocol (L2CAP) Quality of Service (QoS) mechanisms. A brief description of the relevant portions is given below.

Adaptive Frequency Hopping

The high density of radio frequency (RF) systems in the Industrial, Scientific, Medical (ISM) band at 2.4GHz, including microwave ovens, cordless telephones, wireless systems such as the 802.11 b & g, and Bluetooth, has created a need for a standard solution that will enable Bluetooth technology to mitigate other interferences in its working band. This task has been assigned to the coexistence working group of the Bluetooth SIG, which has adopted the AFH as the leading solution.

The AFH is a non-collaborative mechanism, which includes a modified frequency-hopping kernel that adapts itself to the crowded RF medium by dynamically replacing frequencies within the hopping sequence. The channel replacement is based on indications from a channel assessment mechanism, which classifies each of the 79 1MHz channels as "good" or "bad" and makes a decision whether or not to exclude it from the FH sequence. The channel assessment can be based on packet error rate (PER), bit error rate (BER), received signal strength indication (RSSI), or any other applicable metric. The master can update the hopping kernel based on an assessment that it performs itself, or based on channel assessment information that it requests from its slaves. Notwithstanding, the AFH specification does not specify which method should be used for channel assessment, nor does it specify how frequently the assessments should be updated.

The AFH is most effective when the Bluetooth and WLAN radios are not collocated, and it loses its effectiveness if all the band is filled with 802.11 interference. For these reasons AFH has been found less applicable for PCA-based devices during collocated usage scenarios (i.e., when both WLAN and Bluetooth radios have to operate simultaneously). However, AFH is being supported as it forms a very useful feature for other non-collocated usage scenarios. Other methods for enabling coexistence in collocated scenarios are described later.

Extended Synchronous Connection Oriented (eSCO)

The eSCO extends the existing SCO connection by defining a new synchronous logical link, which is more reliable than the legacy SCO. The higher reliability is achieved by applying Cyclic Redundancy Check (CRC), which allows retransmissions of corrupted Extended Voice (EV) packets to each of the eSCO packets. The new mechanism adds three new packet types: EV3, EV4, and EV5, which are defined with regards to the legacy SCO channels as follows:

EV3 = HV3 + CRC

EV4 = 3 slot HV2 + CRC

EV5 = 3 slot HV3 + CRC

The eSCO uses a separate active member address (AM_ADDR) called a logical transport address (LT_ADDR), which enables a separate automatic repeat request (ARQ) scheme to be used in parallel with the existing one that is used for asynchronous connectionless (ACL) transmissions of that specific Bluetooth mechanism. The eSCO can be used for transparent synchronous data as well as for audio. Since eSCO enables synchronous *asymmetric* links (as opposed to SCO), data rates such as 384kbps and 564kbps can be achieved. These rates match the data rates enabled by the 3G Universal Mobile Telecommunications Service (UMTS) systems. This makes eSCO most suitable for audio and video streaming applications over 3G networks. Moreover, the eSCO enables a new way to combat the audio quality degradation that results from coexistence interference in the 2.4GHz band. In addition, since it reduces the collision probability with a transmitted WLAN packet, it increases the throughput of the WLAN system (compared to usage of HV1 or HV3, [5]), so eventually both sides benefit from the usage of the eSCO feature.

Future Improvements of WPAN Radio

Future releases of the Bluetooth specification will introduce two new increased rate enhancements for the basic rate (BR): the medium rate and the high rate.

1. *Medium Rate*–The medium rate mode (MED) [13] provides a data rate enhancement targeted at usage scenarios where a 2x or 3x increase is beneficial. The rate increase is enabled by applying different modulation schemes to only the packet payload, while the rest of the payload remains modulated using the legacy 2-GFSK modulation. The MED enables gross air rates of 2 Mbps using $\pi/4$ Differential Quadrature Phase Shift Keying (DQPSK) modulation, and optionally 3 Mbps using 8 Differential Phase Shift Keying (DPSK) modulation. The achievable user data rates are up to 1.45 Mbps and 2.18 Mbps, respectively. The MED is an evolutionary extension to Bluetooth technology that is designed to operate concurrently with Bluetooth 1.x slaves in the same Bluetooth network (a.k.a. piconet). It reuses most of the existing Bluetooth 1.1 and 1.2 functionality and requires a major change only in the PHY layer. In addition, there will be further minor changes in the Link Manager (LM) and Host Controller Interface (HCI).

2. *High Rate*–The high rate (HR) [11], [12] is a proposal for Bluetooth 2.0 that combines new MAC and PHY layers. The HR aims at achieving data rates that are 10 times higher than the rates enabled by the BR. The new MAC layer introduces the idea of a token-based distributed MAC, supervised by a single device called a "supervisor." The role of the supervisor is mostly to facilitate fair token distribution among the network peer devices. This concept is proposed in order to reduce the relatively high latencies and overhead imposed by the centralized MAC scheme that is used in the 1.x Bluetooth mechanisms.

   The air interface of the HR is the second dramatic change introduced by the new standard. The most important thing to notice is that the HR channel does not frequency hop. Instead, it statically occupies a 4 MHz channel (5.6 MHz@-20dB). The optimal HR RF band is chosen via a dynamic channel selection (DCS) mechanism, which picks up a carrier out of a set of 70 possible, 1 MHz separated, RF carriers. This new air interface introduces an inherent coexistence problem with legacy Bluetooth mechanisms, similar to the one introduced for example, by 802.11. However, the HR specification proposes to tackle this problem by applying the AFH scheme in the legacy FH systems.

   The HR baseband uses an adaptive modulation scheme of M-ary phase shift keying (PSK) differential modulation with a symbol rate of 4 Msymbols/s. Three types of modulation schemes (which can be adapted according to QoS requirements and channel conditions) have been defined as follows:

   - 8-DPSK - 12 Mb/s
   - 4-DPSK - 8 Mb/s
   - 2-DPSK - 4 Mb/s

Finally, Table 1 summaries the main features of the two new radio enhancements of the Bluetooth specification. The table compares the MED and HR in terms of data rates, control, MAC method, Data Link Control (DLC) methods, and Radio Resources Management (RRM) methods.

Table 1: Comparison of medium and high rate

	Medium rate	High rate
user rates (Mb/s)	0.7, 1.4, 2.1	3.8, 7.6, 11.4
control	Centralized	distributed
MAC	Polling	token
DLC	Packet	segment
RRM	FH	DCS

The introduction of the HR might bring up again the debate about the need for Bluetooth technology in parallel with Wi-Fi networks. Nonetheless, the HR should be considered only as a natural evolution of the WPAN, which complements the Wi-Fi technology. The HR is proposed as an optional extension, not a replacement for the Bluetooth 1.x and is designed to live concurrently with the BR piconets. In summary, HR addresses the increasing need for high-speed cable replacement and reliable applications such as high-quality audio and video gaming.

Bluetooth Profile Advancements and Trends

Much of the success of Bluetooth technology in the last couple of years should be credited to its wireless audio applications—the headset and hands-free devices. This trend towards wireless audio devices, powered by the automotive industry and consumer demand, has led to the development of advanced profiles for enhancing the audio quality and increasing the range of applications that can be supported by such devices. New headset designs being released offer, besides a cool shape, an increased talk-time with enhanced audio quality, which are enabled by these hands-free profiles (HFP) and headset profiles (HSP). The dramatic improvement in the audio quality is achieved by incorporating digital signal processors (DSPs) and algorithms for echo cancellation and noise suppression in the headset design, and by using the 1.2 new features such as AFH and eSCO for mitigating local interference.

More and more headset designs are being based on the HFP, which is slowly taking the place of the HSP. The HFP enables not only basic audio connectivity, but also mobile phone remote-control capabilities and is applicable to both headset devices and car-kit hands-free systems. Another profile that extends the HFP features is the phone access profile (PAP), which enables full remote control of the mobile cellular phone through a car phone.

The additional processing power provided by DSPs in the headsets enables advanced profiles such as the advanced audio distribution profile (A2DP), which defines the protocols and procedures that are used to implement distribution of high-quality audio content in mono or stereo (e.g., MP3) on ACL channels. These improvements are geared towards increasing the market penetration of mobile phones enabled with Bluetooth technology.

The Personal Area Networking Profile

The Bluetooth SIG Personal Area Networking (PAN) Working Group (WG) has defined a new protocol and profile to enable existing networking applications to work over Bluetooth links by using existing networking protocols such TCP/IP. The new profile is called the Personal Area Networking (PAN) profile. The Bluetooth SIG defines profiles to specify how devices should use various protocols to ensure interoperability between devices supporting these profiles. In addition to the PAN profile, a specification for a new protocol called Bluetooth Networking Encapsulation Protocol (BNEP) has been developed. The BNEP specification provides an "Ethernet-like" layer in which existing networking applications can function over Bluetooth links without any additional changes.

The PAN profile brings a rich networking environment to devices enabled with Bluetooth technology and explodes the usage of these devices. The PAN profile, along with BNEP protocols, aims to provide an Ethernet-like networking infrastructure for devices enabled with Bluetooth technology. Hence, the PAN profile opens a huge window for Bluetooth products to incorporate TCP/IP-based protocols and applications seamlessly. Essentially, the PAN profile makes the Bluetooth mechanism an integral part of a network that provides the user voice and data on the go.

There are many other profiles that are geared to support printers, audio, and video usage models; they are not addressed or described in this paper.

Ultra Wideband

Ultra Wideband (UWB) is a high-speed, low-power, wireless technology designed for distances less than ten meters. The IEEE 802.15 task group (TG) 3a [6] was formed in early 2003 and is working on defining a specification for a UWB physical layer for the IEEE 802.15.3 protocol. Currently, the IEEE 802.15 TG3a WG is evaluating proposals against their functional requirements criteria. The IEEE 802.15 TG3a requires data rates of 110 Mbps or faster and will have Quality of Service (QoS) support focusing on supporting high-rate multi-media. As the UWB standardization effort progresses so will its importance as a WPAN technology.

Internet Engineering Task Force Mobility Drafts

During the past several years the Internet Engineering Task Force (IETF) has been working on defining protocols to enhance mobility. One protocol in the final stages is Mobile IPv6 [7]. Mobile IPv6 defines a mechanism as something that allows a mobile device to maintain the same IP address as the mobile device moves around the Internet. By maintaining the same IP address, the application will continue to function and maintain all networking connections. This capability allows mobile hosts to move from network to network and switch networking technologies, while maintaining network connections. There are many other IEFT drafts that improve the capabilities of mobile nodes as they move throughout the networks [8].

Java Specification Request - 82

JSR-82 is a Java Specification Request that was issued in 2002 by an expert group chaired by Motorola with additional members such as Rococo Software, Ericsson, Nokia, and others. Its goal was to standardize a set of APIs to allow Java technology-enabled devices to be integrated into a Bluetooth environment.

The JSR-82 specification was defined for the target platform of Java* 2 Micro Edition (J2ME). The J2ME is a subset of the Java 2 platform technologies and is optimized for portable devices such as PDAs and mobile phones. The J2ME is composed of two modules: Configuration and Profiles. The Configuration module provides the minimum set of classes and virtual machine (VM) features that must be present for a particular set of devices. The Profiles module, which is layered on top of the Configuration module, provides the application developers APIs for a particular set of devices. In the context of mobile phones, the Configuration module is called the Connected Limited Device Configuration (CLDC), and the Profile module is called the Mobile Information Device Profile (MIDP). The CLDC provides a fast, small footprint virtual machine (called a KVM) and a stripped-down Java API subset. The MIDP provides a set of user interface components, a persistence mechanism, and a Hypertext Transfer Protocol (HTTP) connection capability.

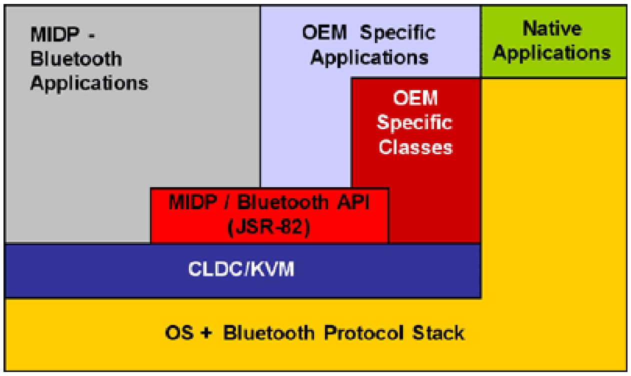
MIDP - Bluetooth Applications
OEM Specific Applications
Native Applications
OEM Specific Classes
MIDP / Bluetooth API (JSR-82)
CLDC/KVM
OS + Bluetooth Protocol Stack

Figure 3: JSR-82: CLDC+MIDP+Bluetooth software architecture

The JSR-82 API provides support for only four fundamental Bluetooth profiles including Generic Access Profile (GAP), Service Discovery Application (SDAP), Serial Port Profile (SPP), and Generic Object Exchange (GOEP), and it has limited access to the stack for Bluetooth radio control. Some basic functionality, such as SCO connection establishment, or resetting the Bluetooth radio, etc., has been excluded and is not being supported by the JSR-82 API. The software architecture of a system enabled with JAVA and Bluetooth technology that supports the JSR-82 is described in Figure 3.

To summarize, the JSR-82 opens for the PCA a new dimension for peer-to-peer networking applications, such as multiplayer wireless games in which users can interact with each other over the Bluetooth link, or interact with a central "data kiosk" for gaming or information-downloading applications.

CHALLENGES AND SOLUTIONS FOR PLATFORM AND HANDSET INTEGRATION

Integrating multiple wireless technologies in Intel Centrino mobile technology and Intel Personal Internet Client Architecture (PCA) creates several challenges. This section of the paper focuses on Wireless Personal Area Network (WPAN) integration for the Intel Centrino mobile technology and PCA platforms.

Bluetooth Integration Considerations in Notebooks

Hardware Integration

Bluetooth specifications define various interfaces that can be integrated with the main system. They are as follows:

- Universal Asynchronous Receiver Transmitter (UART): Bluetooth radio is integrated using a UART connection on the same printed circuit board (PCB) as the main system. This is the preferred method for

*Other brands and names are the property of their respective owners.

most hand-held devices, since it provides cost-effective solutions.

- Universal Serial Bus (USB): Bluetooth radio is integrated as a USB client device. This is a *de facto* standard for mobile PC platforms.
- Card Bus: During the initial stages of Bluetooth development, many vendors provided Bluetooth radio as an add-on card for mobile PC platforms.
- Peripheral Components Interconnect (PCI)-Express: In future mobile PC platforms, one can expect usage of the PCI-Express to be integrated into Bluetooth radio. The PCI-Express offers better power management, software compatibility, and lends itself to integration into the Bluetooth radio on the back lid of a mobile PC platform. Studies have found that integrating a Bluetooth radio on the lid offers superior radio frequency (RF) characteristics.

Power Management Design Considerations for Mobile PC Platforms with an Integrated Bluetooth Radio

Extending battery life is an important objective for the Intel Centrino mobile PC platforms. On Centrino platforms, USB has become a *de facto* standard for integrating a Bluetooth radio on the platform. USB supports multiple techniques of low-power modes. Due to current limitations of the platform, USB selective suspend has become a preferred mechanism for the power management of the Bluetooth radio. In this method, the Bluetooth USB driver suspends USB transactions, when the Bluetooth radio is idle for a few seconds. This allows the PC processor to go into low-power mode, thereby saving battery life significantly.

Wake on the Bluetooth host controller will become one of the key features for future mobile PC platforms. This feature allows a USB-based Bluetooth radio to wake up the mobile PC from the standby mode. The feature provides additional battery savings, while providing the ability to respond to external Bluetooth host controller events. The Bluetooth specification provides a facility for setting up multiple event filters on the Bluetooth radio by the host. Typical event filters are *event* on Bluetooth connection from a specific Bluetooth unit and *event* on paging from any Bluetooth unit. To support Wake on the Bluetooth host controller, the host processor typically sets some event filters before going to standby mode and then goes to standby mode to save battery energy. In this scenario, Bluetooth radio is still powered and wakes the system when it recognizes a particular event that matches the set event filters.

SW Structure—Protocol Stack, Profiles and Applications

Specification of the Bluetooth system defines a layered software structure covering radio baseband systems, link layers, higher level protocols, and application profiles.

Bluetooth radio chip vendors provide baseband support and export the standard HCI interface. Operating system (OS) vendors and application vendors implement the rest of the software stack and applications, which run on the host processor.

In some cases, Bluetooth radio chip vendors are implementing application-level profiles such as Serial Port Profile (SPP) on the chip itself. In this case, the Bluetooth radio acts like a wireless serial port and enables legacy systems without requiring software changes on the host.

An application with Bluetooth wireless technology operates in a wireless environment unlike a typical application that operates in a wired environment. This wireless scenario forces the application to discover other Bluetooth products and services and connects to them securely before exchanging data. The Bluetooth SIG defined a number of profiles to enable service discovery, authentication, and various data transfer mechanisms.

The usage of various profiles can be explained by enumerating a typical conversation between two Bluetooth products.

A Bluetooth product uses a Service Discovery Application profile (SDAP) to discover services offered by nearby willing Bluetooth products. If these Bluetooth products are personal devices, a user can pair them using a Generic Access Profile (GAP) thereby placing significant security barriers for other Bluetooth products to pass. GAPs provide access definitions for establishing basic link-level connections between the devices. The devices can enter into a data transfer phase by using a Serial Port Profile (SPP), which was intended to provide a basic cable-replacement mechanism that turns a Bluetooth product into a wireless serial port. Using this cable-replacement mechanism, the applications can exchange data between the devices.

Current mobile PC platforms with Bluetooth technology support Generic Object Exchange profiles for image transfers, exchanging and synchronizing Personal Information Manager (PIM) information and dial-up network profiles to establish a dial-up connection using hand-held devices enabled with Bluetooth wireless technology.

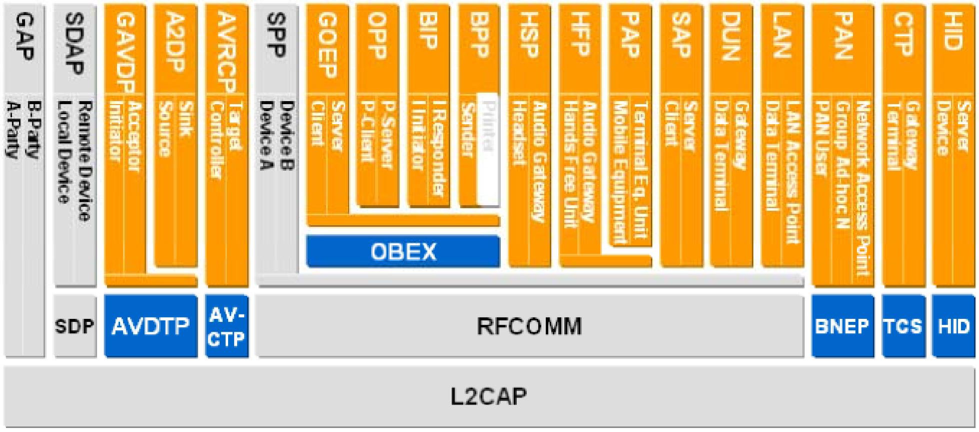

Figure 4: Mobile phone Bluetooth Software Stack and profiles (courtesy Ericsson Technology Licensing AB)

The Bluetooth Human Interface Device (HID) profile addresses usage of wireless keyboards, mice, joysticks, and remote controls based on Bluetooth wireless technology, and it uses the same HID definitions that are defined in USB HID devices. This enables usage of existing HID class drivers of mobile PC platforms, thereby speeding up implementation and adaptation of HID devices that are based on Bluetooth technology.

Further Bluetooth Integration Considerations in Mobile Handsets

The integration of Bluetooth capability into mobile handsets has evolved dramatically since the technology was introduced back in 1998. In its early models, Bluetooth technology was available only as an external dongle or as a sliding card; today, the integration is internal to the printed circuit board (PCB) either as an external single chip, or as an external module comprised of separate baseband and RF chips. In addition, a growing number of cellular and application processors integrate Bluetooth technology baseband capabilities into their silicon.

Today, Bluetooth technology capabilities can be found in a wide variety of mobile handsets and hand-held devices. The main differences between those devices (at least from the Bluetooth wireless technology perspective) are in the set of supported Bluetooth features stemming from the audience for which these phones are targeted. Roughly, we can divide the handsets enabled with Bluetooth wireless technology into four categories: low-end phones, youth phones, feature phones, and smart phones, respectively. Low-end phones usually provide only basic Bluetooth connectivity and therefore support only the headset profile (HSP), some object exchanging capability using the Object Push Profile (OPP) for business card exchanging, and basic printing capabilities using the Basic Printing Profile (BPP). Youth phones interact mainly with multimedia and gaming devices. Their supported profiles therefore focus on multimedia applications such as BIP and A2DP. Feature phones are phones that are intended mainly for professionals and gadget fans. These phones are typically showcases for almost all Bluetooth wireless technology functionality which includes all the profiles seen in Figure 4. And finally, smart phones are a mix of a traditional mobile phone and a PDA. These phones typically support operating systems such as the Microsoft* WinCE* Smartphone* 200x, Symbian*, Linux* and such, and provide Bluetooth wireless technology connectivity that is similar to the one enabled by the feature phones.

The differences in the form factor between mobile computers and hand-held devices impose several design constraints, which may be more relaxed when Bluetooth technology is integrated into mobile computers. The parameters to which special attention should be given include

- *Hardware Interfaces*. When Bluetooth wireless technology is integrated as an external chip (set), the required interfaces include a host controller interface over UART (H4) or a USB (H2), pulse-coded modulation (PCM) for audio. Each of these interfaces consists of four lines. In many cases the UART (H4) is preferable due to the relaxed requirements it imposes on the hardware. It should be noted, however, that UART is not applicable when medium-rate mode (MED) and high rate (HR)

* Other brands and names are the property of their respective owners.

are considered, as it can support only limited data rates, which do not meet the data-rate requirements of the new higher rate radios. Notwithstanding, currently no definitions for new transport layers have been published by the Bluetooth SIG, although USB seems to be a natural candidate.

- *RF Interface*. This issue is mostly applicable when Bluetooth technology baseband IP is integrated into silicon, although it can also be considered as an issue when Bluetooth technology is integrated as an external chipset, composed of a separate baseband and RF chips that come from different vendors (an option which is less popular today). The chip design must then include an interface to an external RF module. Currently, there is no standard interface protocol between the Baseband IC and the RF unit defined in the Bluetooth specification. However, some Bluetooth Working Group members (Ericsson, Nokia Mitel, Intersil, and Philips) have prepared the BlueRF specification for a standard interface between Baseband and RF. An 8-line interface is required for a bidirectional port, while a 14-line interface is required for a unidirectional port.
- *Electrical Compatibility.* When choosing an external Bluetooth host controller, attention should be given to the electrical compatibility of electrical levels, polarity, protocol, and timing between the host and the Bluetooth controller. As an alternative, level-shifting devices may be used for fixing the electrical mismatches, with the drawback of increased Bill-Of-Materials (BOM) and PCB area.
- *Power Management.* Power management presents a major design challenge especially in hand-held device design, due to the smaller battery size that can be considered in such devices, which directly affects the standby and talk time of the device. The UART H4 protocol does not have a data-recovery mechanism, which makes it unsuitable for reliable communication that involves low-power/sleep periods. Data that are sent while the receiving party is asleep are lost. Furthermore, since data are being sent over UART as a stream of bytes, there is no indication as to which of the bytes has been lost. Since the 1.1 release of the Bluetooth specification did not include a standard definition for power management over UART-H4, most commercial Bluetooth chipsets available today rely on proprietary solutions that may vary from one vendor to another. The drawback of these solutions is that they are usually based on additional pins that need to be allocated for acting as wakeup pins for the Bluetooth host and Bluetooth host controller. There is a definition for a new standard (that is not part of the Bluetooth 1.2 specification), called a 3-wire HCI (H5), which, as implied by its name, consists of only 3 lines of the UART interface: TX, RX, and GROUND. The 3-wire HCI implements a protocol stack based on the Serial Line Internet Protocol (SLIP). The SLIP, which is layered on top of the UART driver, is used for transforming the unreliable flow of data bytes over the UART, into a reliable flow of data packets. The H5 defines a data-recovery mechanism as well as a hand-shaking protocol in software that enables bidirectional power management between the host and the Bluetooth host controller over the 3 UART pins. The power-management feature requires that the RX line on each side have the capability of being configured as an asynchronous interrupt line that could wake up the receiving party. In summary, the advantage of H5 for power management is in the saving of the two extra UART pins that are used for flow control (Request To Send (RTS) and Clear To Send (CTS)). The drawback of using it is the requirement for major modifications in the host controller interface (HCI) software and firmware, including the implementation of an additional protocol stack for the UART.

  As mentioned earlier, an alternative solution for the power-management problem is to use the HCI over USB (H2) specification [1], which does include standard interface lines for power management. The main drawback of using USB is the requirement for an implementation of USB host module in silicon, which is considerably more costly and complex than UART H4 silicon implementation.
- *Clocking.* Using a common clock for the entire system (Bluetooth, APPS, COMM, etc.) is desirable for many reasons, including reduced BOM, reduced power consumption, and the ability to control the Bluetooth clock by using a power-management chip. However, in mobile handset Bluetooth applications that involve audio connections such as HSP and HFP, the requirement for a common clock becomes even more critical, as the entire audio path (the audio codec, vocoder, echo canceling, noise suppression, etc.) is synchronized with the cellular base station reference clock. Using a separate clock for the Bluetooth host controller, which is not being tracked and corrected (compared to a VCTCXO clock, for example), would result in an accumulated timing drift error in the sampling time, and this will cause an annoying “clicking” sound to be heard at the near and far ends of the audio link.

CHALLENGES AND SOLUTIONS FOR MOBILE WIRELESS PLATFORMS

There are many challenges for wireless platforms and these challenges are significantly increased with multiple wireless technologies. Theses challenges range from low-level hardware issues such as platform designs to high-level issues that affect applications software. In this section we discuss some of these potential challenges and suggest possible solutions.

Wireless Security

With wireless connectivity, mobile devices are vulnerable to many additional security threats as compared to devices with wired connections. With the use of wireless networks, all data exchange via a wireless connection should be encrypted to prevent others from gaining access to that information. Encryption may be supported in several layers of the device's communication stacks. Both 802.11 and Bluetooth support encryption over the wireless link and should be used. Additionally, wireless devices, which connect up to public networks, should use Virtual Private Networks (VPNs) when connected to those networks. VPN encrypts communications from the source to the destination to prevent others from gaining access to that information as it travels the Internet.

Moreover, wireless mobile devices are vulnerable to attacks in which others may try to gain access to the mobile devices or effect their operation. Due to this threat, mobile devices should implement firewalls and other mechanisms to limit unauthorized access to the mobile devices.

Roaming

As wireless devices move from place to place, they change the location of network connectivity. As this happens the Internet Protocol (IP) address may change. The IP address is used by the Internet to deliver network communication to the correct device. Applications use IP addresses to share information across the Internet. If the IP address changes, then applications may be prevented from sharing information. One possible method to solve this issue is to use Mobile IP. With Mobile IP, the mobile device uses its home IP address for its home network as the destination for all communications. When the mobile device is at its home network, it receives all of the communications directly. When the mobile device moves away from its home network, it gets a new IP address from the visiting network and tells a device called the Home Agent (HA) what the new visiting IP address is. The Home Agent redirects network communication from the mobile devices home network to the new visiting IP address. This allows applications on other devices connected to the Internet to always use the mobile device's home IP address in order to communicate with that device. This allows users and their applications to maintain network communications as the mobile device moves from network to network.

Quality of Service (QoS)

Wireless networks are typically shared with all of the users in an area. That area depends on the technology; for example, a typical range for networks is 10 meters (m) for WPANs, 100m for WLANs, and 1000m for WWANs. Due to the sharing of the network, users may not get enough network bandwidth for their applications. Many of today's network bandwidth is shared without any guarantees. QoS mechanisms allow devices to request a guaranteed amount of access to the network. QoS mechanisms guarantee different network parameters such as data rate and latencies. By having QoS, applications can request the QoS parameters needed for their application and have a good user experience.

Bluetooth Wireless Technology Coexistence

As described in the introductory sections, the technologies that can enable the dream of "anywhere, anytime connectivity" impose the integration of cellular/Wireless Wide Area Networking (WWAN) such as GPRS, EGPRS or W-CDMA, Bluetooth technology, 802.11 WLAN, and GPS on a single platform. Side by side with the advantages that are enabled by an integration of multiple wireless technologies, such as greatly enhancing connectivity to the enterprise Intranet and the Internet and to peripherals such as PDAs, printers, and headsets, simultaneous operation of collocated radios in hand-held or mobile computer platforms presents significant design challenges. The main challenges are as follows:

- Spectral overlapping: The operation of 802.11b/g and Bluetooth technology in the same ISM band causes inevitable transmission collisions and results in QoS and throughput degradation.
- TX Out of Band: This is noise generated at the TX band of one system that increases the RX noise floor of the other system.
- Blocking: This phenomenon is caused by a strong TX signal that overloads the RX front-end.
- RX phase noise: This noise is generated by the local oscillator causing a mixing of TX signals into the RX band.
- Inter-modulation Interference: Inter modulation is caused by the collocated existence of multiple radios generating undesired inter-modulation frequency products that need to be filtered.

- Frequency planning: Careful frequency planning will ensure that frequency source harmonics don't fall in the wrong bands.
- Small factor platforms: The small form factor (especially the one that is considered in mobile phones and hand-held devices) means that several radios may share the same antenna, which can result in limited isolation between antennas.

The design challenges mentioned above are true for both mobile notebook computers as well as for hand-held devices or mobile phones. Hence, the following methods for handling and eliminating undesired RF effects apply to both designs.

Mitigating the effects of TX Out-of-Band noise, blocking, or RX phase noise is usually accomplished by adding filtering at the TX side of each of the collocated RF systems. The price for additional filtering is increased BOM and PCB size, and the insertion loss. In addition, the filtering influences the noise figure (NF) and thus the overall performance. Therefore, careful design tradeoffs have to be made.

Inter-modulation effects can be handled by using more linear devices (at the expense of higher current consumption and cost), or by using additional filtering in the RX side, which affects cost, size, and insertion loss. Using common frequency sources (if possible), additional shielding (which increases the size and cost), or robust conversion schemes, such as direct conversion radio (DCR), may offset the undesired effects of frequency planning.

Up to this point we have reviewed the design guidelines, which are common to both platforms. Yet, there is a difference in the solution approaches between hand-held and mobile computer designs. The main issues that should be taken into account are the form factor and the usage models.

Firstly, the form factor affects the antenna design. A small form factor implies that the antennas will be closer together, which results in limited antenna isolation. In addition, Bluetooth and WLAN radios may share the same antenna, which requires a time-sharing mechanism between the systems as well as tighter design of RF switches and RF terminators to prevent TX energy from leaking into the other system's receiver.

Secondly, the form factor limits the size of the batteries that can be used in the design, thereby implying tighter power-management designs for hand-held devices, as well as tighter requirements for the WLAN design.

The usage models that are considered for mobile notebook computers and mobile phones are also different. As an example, consider a phone call using a Bluetooth headset while data are being streamed over a WLAN connection. Since the audio connections over the Bluetooth headset use reserved time slots for transmission, simultaneous WLAN transmissions in those epochs can greatly reduce the Bluetooth audio quality. The usage of GPS is another example that is unique to hand-held devices; it doesn't apply to mobile notebook computers.

In the next sections we discuss solutions for coexistence that are unique to each of the platforms, separately.

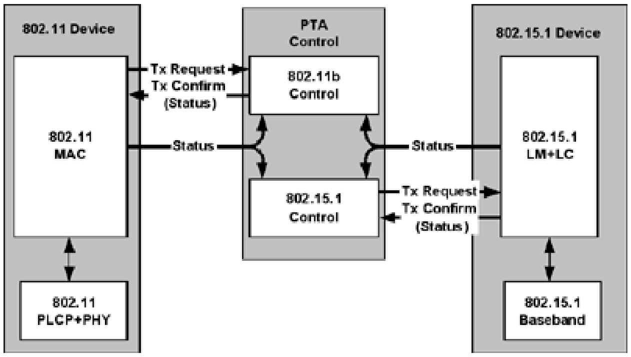

Figure 5: PTA collaborative mechanism for coexistence of 802.11b and Bluetooth technology

Coexistence Solutions in Mobile Handset Design

The integration of GPS in mobile handsets even without the presence of Bluetooth technology and WLAN is challenging, as the design of GPS receivers relies heavily on received noise levels. Higher noise levels require longer detection periods. Furthermore, GPS receivers are not designed to ignore noise bursts. Most of the existing GPS designs did not consider noise due to collocated RF transmitters. Handling such interference requires a change in the design of the receiver: tighter RF design and longer detection periods will be required. In addition, some GPS solutions will not work with a GSM interference of more than -25 dBm, and the performance of most of them starts to degrade at -35 dBm.

The coexistence of Bluetooth technology and WLAN in the same device is perhaps the most complicated of all the coexistence combinations that can be considered inside a hand-held device. Once the requirement for Bluetooth-802.11 coexistence was raised, task forces inside the Bluetooth SIG and the IEEE (IEEE 802.15.2) have been formed to propose standard solutions for the problem. Several collaborative and non-collaborative mechanisms have been proposed for enabling coexistence between 802.11b/g and Bluetooth technology [4], [5]. As mentioned previously, non-collaborative mechanisms, such as the AFH, have been mainly proposed in order to mitigate interferences in the 2.4 GHz range, rather than enabling coexistence in a collocated scenario. Using AFH solely in a hand-held device will not enable handling a voice conversation using a headset enabled with Bluetooth technology with the required sound quality, while 802.11 is trying to upload a packet. These kinds of scenarios that combine audio and data are the essence of the PCA architecture. This requirement impinges on the usage of collaborative techniques, such as Packet Traffic Arbitration (PTA), described in Figure 5, for enabling 802.11 and Bluetooth technology to coexist. PTA uses a “control entity” with the ability to control both the 802.11 and Bluetooth MACs. The control entity acts as a “traffic cop” and implements a hand-shake mechanism with both MACs in order to authorize transmissions in a time multiplexing manner. The decision algorithm gives priority to either Bluetooth or 802.11 according to the transmission type: synchronous audio transmission (SCO or eSCO) may be given priority over 802.11 data traffic. Other data packets may be permitted on a Quality of Service (QoS) basis or on similar decision mechanisms.

The implementation of the controlling entity can be done in many different ways and depends on the specific design. In an external integration of a Bluetooth host controller and WLAN, one option is to use discrete control lines to connect the MACs of the Bluetooth controller and the WLAN to each other. This can usually be implemented using two lines. The air manager can be implemented on either the Bluetooth controller or the WLAN, depending on which of the two has spare processing power for the job. Another option is to connect both MACs to the host (e.g., APPS processor) and implement the air manager inside the host processor. The drawback to this is an increased latency (due to the OS response time to hardware (HW) interrupts) and an increase in the host processor pinout (or in other words, less general-purpose input/output (GPIO) pins for other applications).

Due to the above reasons, a conspicuous market trend points towards collaboration between Bluetooth technology vendors (e.g., Silicon Wave) and WLAN vendors in order to provide a complete integrated solution for Bluetooth-WLAN coexistence. These solutions focus on providing not only a coexistence solution, but also an efficient WLAN design, from the power-consumption perspective, due to the strict power requirements for the design of hand-held devices. Nonetheless, even though those solutions require an air manager to be implemented into the design of Bluetooth-WLAN products, some derivative of an air manager implemented in the host may be required for enabling coexistence with GPS as well.

Radio Frequency Challenges in Mobile Notebook PCs

Antenna designs for notebook PCs are challenging, as generally these kinds of platforms consist of many metallic components such as framing structures, hard drives, and displays. Those can greatly distort the antenna radiation patterns, thereby causing potentially significant RF performance variations, as the notebook PC’s physical orientation is varied relative to the location of the intended communicating device such as an access point (AP), a peripheral such as a printer, and even another notebook PC.

Yet, the biggest coexistence design challenge which the system designers are facing is the enabling of coexistence between the technologies of Bluetooth and 802.11. The Intel Centrino mobile technology includes the Intel Wireless Coexistence System (Intel WCS), which significantly mitigates the interference between 802.11b and Bluetooth technologies.

Intel Wireless Coexistence System

While antenna isolation somewhat mitigates interference between 802.11b and Bluetooth radios, performance is still impacted to some degree. For example, 802.11b data throughput is degraded by Bluetooth interference, even with 40dB of antenna isolation. To further mitigate interference between the Bluetooth radio and 802.11b, Intel WCS [18] was developed as one of the Intel Centrino-enabling technologies. Intel WCS, described in Figure 6, consists of a combination of antenna isolation techniques, a channel exchange, and priority signaling between an Intel PRO/Wireless Network Connection 802.11 solution and a third-party Bluetooth module. Phase 1 of Intel WCS has been implemented: it mitigates Bluetooth interference and restores 802.11b data throughput nearly completely.

Intel WCS is designed to complement the Adaptive Frequency Hopping (AFH) interference mitigation algorithm being developed by the Bluetooth Special Interest Group (SIG). AFH will mitigate the impact of 802.11b on Bluetooth data throughput, but only between AFH-compliant Bluetooth products.

Phase 2 of the Intel WCS will add Bluetooth priority signaling from the Bluetooth module to the Intel PRO/Wireless network connection, resulting in a restoration of connection reliability for both AFH and non-AFH (legacy) Bluetooth modules.

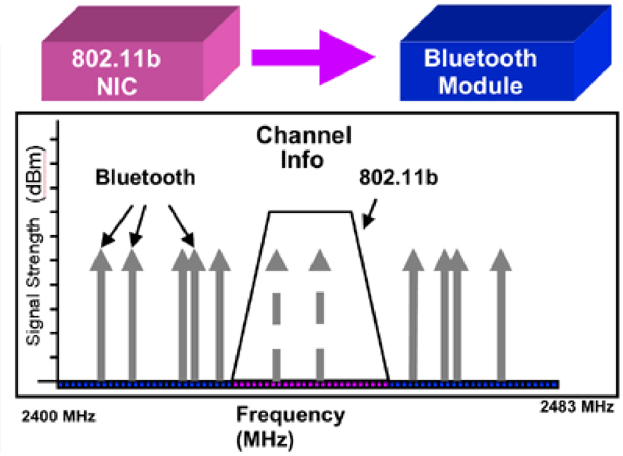

Figure 6: Intel WCS channel exchange

Multiple third-party Bluetooth silicon and module vendors have been enabled to be compatible with Intel WCS, including SiliconWave and Cambridge Silicon Radio. Extensive verification and validation testing has been completed with these silicon vendors, providing a pre-validated system solution to the customer.

In summary, integrating multiple RF technologies into mobile notebook PCs provides new challenges to systems designers, including antenna gain uniformity and interference mitigation. Intel WCS, an Intel Centrino mobile technology, provides powerful interference mitigation between 802.11b and Bluetooth radios and enhances AFH.

SUMMARY

Intel Centrino mobile technology and Personal Internet Client Architecture (PCA) are advanced communication platforms with multiple wired and wireless technologies. With these rich communication technologies, these platforms can be used for many new usage models. Many of these new usage models involve the interaction of Intel Centrino mobile technology and the Intel PCA. Additionally, some of these new usage models are enabled with the use of both the Wireless Local Area Network (WLAN) and the Wireless Personal Area Network (WPAN) wireless technologies. Many new technology innovations have enabled these usage models, and the innovations have been achieved by the use of industry standards to ensure interoperability. For multiple wireless technologies, we discussed in detail suggested methods for integrating these technologies into the platform and talked about how to deal with coexistence issues.

ACKNOWLEDGMENTS

The authors thank Dudi Ben-Eli from CRRL\PCG, Yerachmiel Yeshayahu from SDC\PCG, and Hani Elgebaly from MPG for their helpful comments.

REFERENCES

[1] Specification of the Bluetooth System, Vol. 1, Core, Rev. 1.1, Bluetooth SIG, February 22, 2001.

[2] Specification of the Bluetooth System, Vol. 2, Profiles, Rev. 1.1, Bluetooth SIG, February 22, 2001.

[3] Specification of the Bluetooth System Core Package, Version 1.2, Draft [D4].

[4] *Wi-Fi* (802.11b) and Bluetooth*: An Examination of Coexistence Approaches*, Mobilian Corporation, January 2001, available at http://www.mobilian.com/images/coexistence_of_802_11b_and_bluetooth1.pdf

[5] IEEE P802.15.2/Draft #06, October 8, 2002.

[6] http://www.ieee802.org/15/pub/TG3a.html

[7] http://www.ietf.org/html.charters/mobileip-charter.html

[8] http://www.ietf.org/html.charters/manet-charter.html

[9] http://www.wireless-designer.com/bluetooth_market.html

[10] Jennifer Bray and Charles F. Sturman, *Bluetooth 1.1 - Connect Without Cables*, 2nd Ed., Prentice Hall, Upper Saddle River, NJ 07458, 2002.

[11] Jaap Haartsen, "HIGH-RATE MODE: Getting up to speed," Pre-Congress Bluetooth SIG Associates Seminar, Amsterdam, June 11, 2002.

[12] Bluetooth SIG Radio Working Group, Bluetooth 2.0 Baseband Draft Specification, February 19, 2002.

[13] Bluetooth SIG Radio Working Group, Radio 1.0 Improvements: Medium Rate Baseband Specification Proposal for version 0.7 Team BT1-MED, April 4, 2003.

[14] Java* APIs for Bluetooth* Wireless Technology (JSR-82) Specification, Version 1.0a, Java* 2 Platform, Micro Edition, Motorola, April 2002.

[15] The Intel® Personal Internet Client Architecture – white paper, Intel Corporation, September 2001.

[16] http://www.intel.com/pca/developernetwork/index.htm?iid=sr+pca&

[17] http://www.wireless-designer.com/bluetooth_market.html

[18] G. Chinn, *et al.*, "Mobile PC Platforms Enabled with Intel® Centrino™ Mobile Technology," *Intel Technology Journal, Vol. 7, Issue 2, 2003.*

[19] Bob O'Hara and Al Petrick, *802.11 Handbook: A Designer's Companion*, IEEE Press, 1999.

[20] T. S. Rappaport, *Wireless Communications Principles and Practice*, 2nd edition, Prentice Hall, Upper Saddle River NJ, 2002.

[21] Jim Geier, "Making the Choice: 802.11a or 802.11g," April 15, 2002, available at http://www.80211-planet.com/tutorials/print.php/1009431

AUTHORS' BIOGRAPHIES

Ofer Bar-Shalom is a system engineer in the PCA Components Group (PCG) in WCCG. He has been with Intel for more than five years, and his current focus is on the integration of Bluetooth technology into the various cellular processors' projects in PCG, including GSM/GPRS and E-GPRS. He holds an M.Sc. degree in Electrical Engineering and a B.Sc. in Mechanical Engineering, both from Tel-Aviv University, Tel-Aviv, Israel, and is currently also a Ph.D. candidate at Electrical-Engineering-Systems department in the faculty of Engineering at Tel-Aviv University. His e-mail is Ofer.Bar-Shalom@intel.com

Gordon Chinn is the Radio Frequency (RF) Systems Architecture manager in Intel's Mobile Platforms Group (MPG) and is responsible for enabling mobile platform PCs with multiple, concurrent, embedded radio technologies. He has over 30 years of experience in various RF communications development positions in the commercial and aerospace industry. Gordon holds a B.S. and M.S. degree in Electrical Engineering and Computer Science from the University of California at Berkeley. His e-mail is Gordon.Chinn @Intel.com

Kris Fleming is a senior software engineer for mobile communication architecture in Intel's Mobile Platforms Group. He focuses on various mobile PC wireless computing projects and for the past several years has been the chair of the Personal Area Networking Bluetooth SIG working group and Intel's Bluetooth Architecture Review Board. Kris earned an M.Sc. and a B.Sc. degree in Electrical and Computer Engineering from the University of Maine. In addition, he has an interdisciplinary M.Sc. degree in Computer Science & Computer Engineering from Carnegie Mellon University. His e-mail is Kris.D.Fleming@Intel.com

Uma Gadamsetty is a staff software engineer for mobile communication architecture in Intel's Mobile Platforms Group (MPG). He has been with Intel for more than seven years and his current focus is on various mobile PC wireless communication projects. Uma holds an M.S. degree in Electrical Engineering. His e-mail is Uma.M.Gadamsetty @intel.com

Dynamic Wired and Wireless Networks on Demand

Amber Sistla, Network Architecture Labs, Intel Corporation
Jeremy Rover, Network Architecture Labs, Intel Corporation
Asha Keddy, Network Architecture Labs, Intel Corporation

Index words: dynamic network configuration, network architecture, pseudo-random scenario generation, automation, testing, wireless, validation, roaming

ABSTRACT

The complexity of use scenarios, user device types, and heterogeneity of networks poses an increasing validation and network management challenge. Current networks combine a wide variety of technologies, protocols, platforms, and device types which, when combined with usage models, leads to a combinational explosion. This paper describes the Dynamic Networks on Demand (DND) framework that automates network creation for wired and wireless network topologies and executes resource transitions (e.g., moving a laptop between subnets) all without ever manually plugging or unplugging a wire or configuring a network device. The uses include automated management of local and remote networks, automated demos for promotion of wireless technology, and automated validation of roaming concepts. These combined pioneer mechanisms enable the DND framework to step up to the challenge of provisioning increasingly varied and dynamic networks.

The paper also describes details of the DND framework. The DND framework is used to create hundreds of networks from limited resources and to execute thousands of pseudo-random test scenarios in each network. Using simple human-readable network requests, the DND framework creates a new network through a series of synchronized steps to reconfigure all portions of the new network: Dynamic Host Control Protocol (DHCP) servers, routers, and logical connections. Resource transitions occur by using the same human-readable network requests. The framework also provides error reporting and logging capabilities. The DND framework uses a single control point so the network changes can be controlled remotely. The framework allows for remote configuration of hardware resources as fast as the hardware can handle. The DND framework is designed to be modular and extensible, and can be used with varied devices from different vendors with minimal changes. The paper also highlights the exponential cost reduction due to the reduction of manual intervention (the number of tests as well as the network reconfigurations would not be feasible if done manually) and to the reuse of limited resources, achieved through intelligent and innovative automation.

INTRODUCTION

The introduction of mobile technologies such as Intel® Centrino™ mobile technology fuels network evolution at unprecedented rates, thereby fostering new usage models. The demanding consumers of today need to stay connected anywhere, anytime with adequate security models and ease of use. These demands in turn act as catalysts to combine different infrastructures, topologies, technologies, protocols, and network components. Network components in turn include routers, Dynamic Host Control Protocol (DHCP) servers, virtual private networks (VPN), access points (APs), and user stations (desktops, laptops) connecting over wired and wireless (802.11, Global System for Mobile Communications) protocols with services such as single bill roaming through different wireless hotspots.

In addition, each network component requires different configurations depending on the needs of a particular network. User stations and wireless hotspots add more complexity with variable platforms (e.g., Intel Centrino mobile technology), operating systems, network adapters, manufacturers, etc. Network components can also communicate over different protocols (IPv4, IPv6, Mobile IPv4). The wide variety of components and configurations causes a combinational explosion, with heterogeneous networks of many different types being deployed and used.

Intel Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Different combinations of configurations are possible for initial network setup, and the network must continually adapt as changing needs require the addition or removal of network components in any portion of the network. In addition, networks are also expected to be self-aware and resilient as they adjust to dynamically changing resources on the network.

The manual configuration and reconfiguration to address the many network needs can cause significant delays in network turnaround time and network downtime. In-depth and specific knowledge about each network technology and component is also necessary in order to deploy and manage such networks.

An ever-increasing network management and validation challenge results from the dynamic, heterogeneous, and knowledge-intensive nature of real-world networks.

The Switched Roaming[1] technology (also known as "adapter switching"), developed as a part of the Intel Centrino mobile technology effort, required a validation and debugging environment to tackle the dynamic nature of the technology being developed. Performance, mobility, reliability, and the quality of service the user experiences need to be validated.

The Dynamic Networks on Demand (DND) architecture addresses these challenges by streamlining the network setup process to reduce network turnaround time and by making the technology accessible to both experts and non-experts, allowing real-world networks to quickly be created "on demand."

DYNAMIC NETWORKS ON DEMAND FRAMEWORK OVERVIEW

The Dynamic Networks on Demand (DND) framework is responsible for three unique functions:

1. To create dynamic networks and topologies for known, predefined configurations.
2. To reconfigure networks based on the possible device transitions.
3. To execute network scenarios, perform verification, and report the outcomes from a single control point.

A network scenario for the framework can be defined as these three functions in sequence. The framework creates a network, transitions through various network states, validates the network throughout various stages of transition, then completes the scenario.

To address the three areas of functionality, the DND framework is abstracted into the following four distinct layers: the Control Layer, the Network Management Layer, the Verification and Validation Layer, and the Physical Network Layer (Figure 1). The expandable nature of the architecture allows the addition of components and/or layers at additional levels of the architecture to increase the breadth of functionality.

Control Layer

The Control Layer abstracts the inner workings of the framework and allows users to run generated network configurations. Using network scenarios generated by the Control Layer, network creations and transitions can become a secondary concern to the user.

Users can generate scenarios either randomly or based on predefined network configurations. This layer manages these configurations as well as subsets and expansions of network configurations. Randomization features are available in the context of network scenario creation. Each series of network scenarios can be reproduced by supplying a seed logged in past scenarios.

Based on interactions with the Network Management Layer and the Verification and Validation Layer, the Control Layer determines the current physical layout and state of the network. It accesses the physical network configuration through the Network Management Layer to perform network scenarios and transitions on the physical devices of the network. When the physical devices are added or removed from the pool of available network devices, the available device configuration file generally requires only one update to dynamically propagate changes through the entire framework at run-time.

When interactions with the other layers indicate network failures (e.g., device fails to obtain Internet Protocol (IP) address, network is saturated), the DND framework provides mechanisms for performing graceful state recovery by detecting errors from its interactions with other layers to log and restore proper network functionality.

The Control Layer exists as the single control point for the framework. In addition to synchronizing the actions of the layers, another advantage of a single control point is remote management. The functionality of the framework can be accessed from one point through a console directly on a machine in close physical proximity to the network or via a remote login session (e.g., telnet).

[1] Switched Roaming technology: The process of selecting one network interface in the presence of multiple interfaces based on user-defined preferences. This technology is used with today's network cards to move through network environments without restarting applications.

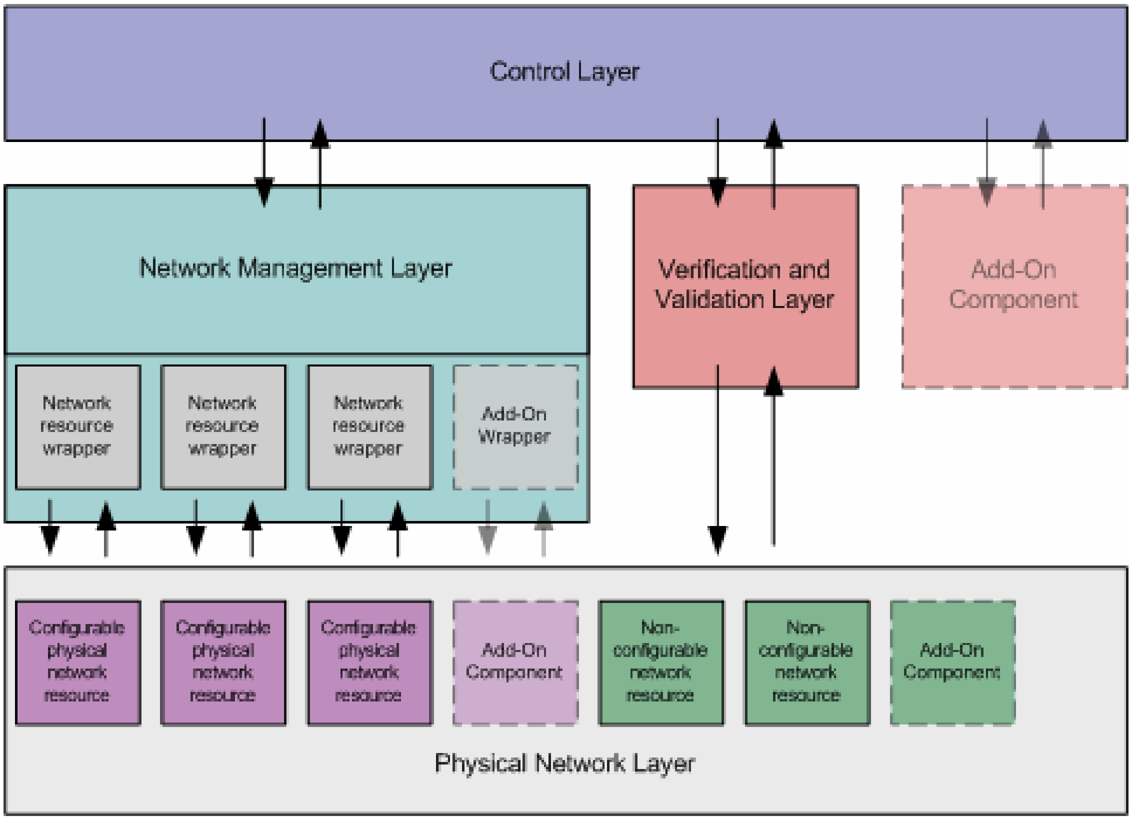
Control Layer
Network Management Layer
Verification and Validation Layer
Add-On Component
Network resource wrapper
Network resource wrapper
Network resource wrapper
Add-On Wrapper
Configurable physical network resource
Configurable physical network resource
Configurable physical network resource
Add-On Component
Non-configurable network resource
Non-configurable network resource
Add-On Component
Physical Network Layer

Figure 1: Layer architecture

Network Management Layer

The Network Management Layer is responsible for all network configuration changes and for maintaining the current physical network state information. This layer was built with reuse and extensibility in mind. All other layers extract network state information from the Network Management Layer.

The Network Management Layer is also a stand-alone component for managing and reconfiguring network devices. This layer can directly create networks and perform transitions when debugging or when networking validation requires manual interaction with devices. In addition, network creation initializes the network state of all devices.

Dynamic network creation does more than just lay out network devices into network topologies. The framework also manages the IP address allocations for all hardware devices on the network and generates a readable text file that reports each IP address to facilitate communication across the network. After every network creation or transition, the Network Management Layer generates an accurate snapshot of the current network state.

The current network state snapshot file also enables the transitioning of devices through a network. In addition to containing the dynamic IP addresses information, the file also includes a list of possible network transitions that a device can perform based on the device capabilities and the configured network infrastructure (e.g., if a device contains an 802.11a network adapter and an 802.11a access point (AP) exists on the network, the device would be able to make a wireless transition to 802.11a). After a transition, generated transition lists are updated for all devices in the network since some network transitions affect multiple mobile devices.

The framework can also be extended by writing wrappers for any additional required configurable devices such as Dynamic Host Control Protocol (DHCP) servers, routers, etc. Adding non-configurable devices such as hubs, laptops, etc., requires minimal effort. By keeping track of all available components in the network as well as the current state of the network, this layer provides flexibility to adapt to changing network needs and resources.

Verification and Validation Layer

The Verification and Validation Layer abstracts all devices used to validate the current network configuration. These devices include packet sniffers, traffic generators, and other network validation devices. Third-party verification tools can be added to this layer to provide seamless accessibility to complex network analysis and traffic generation tools like Chariot*, IXIA* and Shomiti*.

A significant feature of the DND framework is the ability to make the validation devices mobile. Network scenario creation puts expensive network validation tools on the particular subnet that requires validation. In traditional networks, dynamically allocating validation devices means running scenarios over and over while manually moving the validation devices between subnets of the network to capture pertinent details of the network operation. The DND framework makes moving hardware just as easy as requesting the hardware at network creation time.

This layer extracts device information such as the dynamically changing IP addresses from the Network Management Layer. The Verification and Validation Layer does not retain network-specific information throughout scenarios.

This layer reports all errors and logs to the Control Layer. At scenario completion, Verification and Validation errors and logs enable the Control Layer to make judgments about whether to continue, to generate a new scenario, and/or to recover from error conditions. In this way, the detection of network errors triggers graceful network state recovery in the Control Layer.

Physical Network Layer

The Physical Network Layer contains all the physical network devices. Adding physical network devices to this layer requires collaboration with the Network Management Layer.

Configuring traditional networks requires in-depth knowledge of each network component's configuration. The DND framework, on the other hand, abstracts all network devices to a resource name. If certain network device functionality is required for a network scenario, the Control Layer requests that functionality in a network request.

The ability to move network components anywhere on a network is useful for reactive network devices such as clients and for connectivity devices such as APs. However, there is often a need to reconfigure components based on their position and the particular network scenario being executed. The DND framework addresses this necessity with network resource wrappers in the Network Management Layer.

A network resource wrapper is required to configure network components. Objects such as routers, Virtual Local Area Network (VLAN) switches, and DHCP servers require reconfiguration and therefore have network resource wrappers located in the Network Management Layer. Each network resource wrapper maps directly to a configurable physical network resource. When new configurable physical network resources are added, accompanying network resource wrappers are written to abstract the functionality for the Network Management Layer.

Non-configurable network resources can be added and removed from the network at will, as no additional framework support is required to interact with them.

Layer Interactions

The interaction of the layers is shown in Figure 2.

A user activates the Control Layer (not shown). The network scenario then starts with the Control Layer querying the Network Management Layer's network state to determine if the scenario can be performed given the current network configuration.

If the network scenario is supported in the current network state, the sequence skips to the Verification and Validation Layer (described in the following paragraph). If the network scenario is not supported in the current network state, the Control Layer resolves the network scenario into a network configuration and creates a simple, readable text file containing a network request. The network request file contains one or more subnets as well as the starting position on the network for mobile nodes[2]. Using the network request from the Control Layer, the Network Management Layer creates a new network and reports success or failure back to the Control Layer.

Barring failures, the Control Layer triggers the Verification and Validation Layer, which performs any

[*] Other brands and names are the property of their respective owners.

[2] Computing devices with one or more network connections.

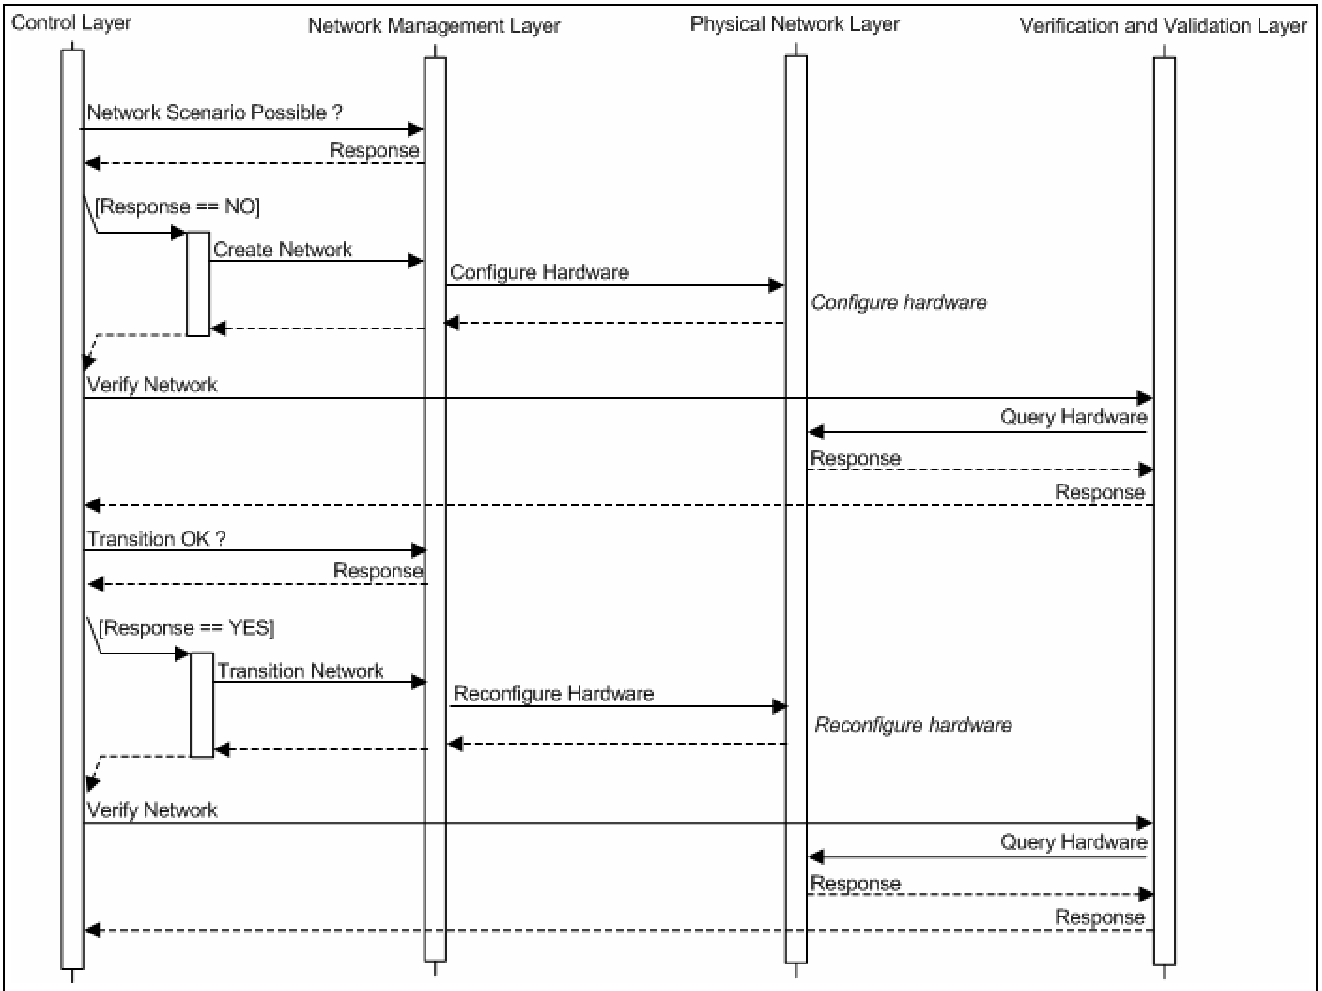

Figure 2: Network scenario sequence diagram

network verification and/or validation tests and reports the findings back to the Control Layer.

If required, the Control Layer queries the Network Management Layer about whether a network transition of a particular network device is supported by the current network configuration. The Control Layer then requests the second half of the network scenario.

For an unsuccessful transition, the Control Layer reports the error and progresses to the end of the sequence as described in Figure 2. For a successful transition, the Control Layer again prompts the Verification and Validation Layer to perform network verification and/or validation and report the findings back to the Control Layer.

At this point the sequence can be iterated, enabling multiple network scenarios to occur back-to-back. The sequence can terminate and report findings for the just-completed network scenario, complete with a pseudo-random seed to reproduce the scenario if required.

REFERENCE IMPLEMENTATION FOR VALIDATION

Mobile technologies introduce new challenges in validation environments, and the Dynamic Networks on Demand (DND) framework addresses some of these problems.

The implementation of the DND framework addresses validating the mobility features of the Intel Centrino mobile technology platform. The platform ships with multiple network interfaces and exposes users to problems such as deciding which network connection to use as the primary connection. The Switched Roaming technology developed in conjunction with the Intel

Centrino mobile technology effort addresses this problem. The Switched Roaming technology selects which network connection to use when multiple connections are present. The DND framework creates network scenarios where multiple network connections are present and transitions a client throughout a network.

The Switched Roaming technology component includes virtual private networks (VPN) auto-launching functionality. When a client obtains a network connection with a network that requires a VPN, the component automatically launches this software. The DND framework creates network scenarios with VPNs to validate functionality.

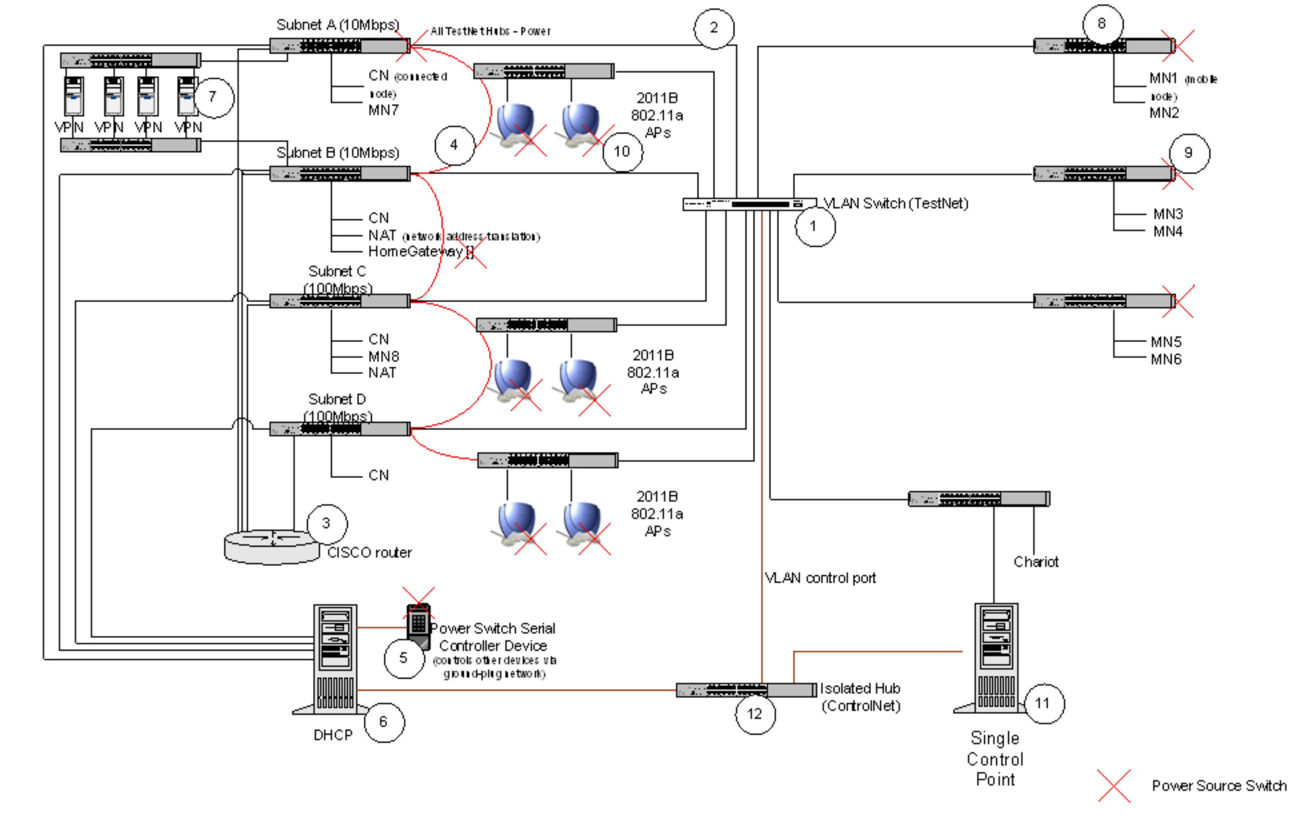

Figure 3: Network diagram

The implementation performs multiple functions. Listed below are some of these functions, the corresponding architectural layer responsible for performing the function, and references to the implementation in Figure 3.

1. **Creating networks**: The Network Management Layer is responsible for creating networks. To create a network the following actions must occur:

- *Determine whether the network can be created.* The Network Management Layer must determine the feasibility of a network scenario before it is created. Only valid scenarios will result in a configured network.

- *VLAN switch reconfiguration.* Configuring the VLAN switch (Figure 3, #1) requires a network resource wrapper. The VLAN switch allows the framework to group hubs together programmatically into logical subnets. Also, ports on the switch that contain router interfaces, DHCP interfaces, and other network devices must be placed in the same VLAN to create an addressable, routable, and valid subnet.

In the example, the router interface, DHCP interface and subnet interface are all located on one hub and plugged directly into the VLAN (Figure 3, #2). The VLAN switch then creates VLANs to join subnets on the left side of Figure 3 to one or more hub(s) on the right side of Figure 3.

- *Create routable subnets.* A router requires a network resource wrapper for configuration (Figure 3, #3). Since the DHCP server communicates the default gateway of each network subnet, the router must have its routing interfaces configured to correspond to the IP addresses that the DHCP server leases to clients. The Network Management Layer creates the associations and configures the routable subnets.

- *Force clients to renew their IP addresses to realize the new network.* The Network Management Layer uses a programmable power source switch to enable this functionality. The DHCP clients on the network nodes respond to their network link being up or down. When a user unplugs the network cable and then plugs it into another network, the DHCP client realizes the change and acquires a new IP address. To force this same response from the DHCP client, the same message to the DHCP client can be produced by powering off a network node's hub, making the client assume it is unplugged from the network. A power source switch located on the hubs of all DHCP clients creates this response (Figure 3, #4). The Network Management Layer controls the power state of these devices via a serial connection to the power source switch controller device (Figure 3, #5)

- *Obtain a valid and predictable IP address.* The DHCP server's network resource wrapper allows IP addresses to be specified for each client on the network in every subnet configuration. The Network Management Layer maintains the IP address list and displays only the current IP(s) of the client in any configuration of the network (Figure 3, #6).

- *Support VPN network configurations.* The main challenge of supporting VPNs is that a number of VPNs require static IP address configurations. Special consideration is given to VPNs in the creation of a network due to their static configuration. If a network requires VPNs, the Network Management Layer creates the specific subnets required for communicating with the statically configured VPN. This is accomplished by configuring the predictable IP addresses on the DHCP server with the required subnet IP addresses that correspond to the specific VPN. It also requires that the router be configured to isolate network traffic on either side of the VPN so that only VPN traffic is routed (Figure 3, #7).

- *Add/remove network resource.* Adding new devices to test in the scenario and/or removing them throughout the lifetime of a project require the ability to quickly update the current availability of resources. The DND framework allows devices to be easily added and removed from hubs (Figure 3, #8). The Network Management Layer queries the hardware list at run-time in order to use the currently available hardware.

2. **Moving network resources after creation (Roam):** The Network Management Layer is responsible for moving resources after a network is created, which creates its own challenges to overcome. Among these challenges are the following:

   - *Change subnets.* Changing which subnet a network resource is located on requires VLAN changes. Changing the port associations of the VLAN allows the network resource to be moved to other subnets.

   - *Obtain a new valid and predictable IP address.* After the VLAN is changed, the network resource will still have its old IP address from the previous subnet. In order to change the IP address, the DHCP client must realize the network connection change. A power-source switch located on the hubs of all DHCP clients creates this response (Figure 3, #9). The Network Management Layer controls the power state of these devices via a serial connection to the power source switch controller device (Figure 3, #5). After the power is restored, the DHCP client picks up the new IP address to complete the move.

   - *Change network topologies (Wired/Wireless) and allow overlapping network interfaces.* Changing which network topology to which network resources are connected requires additional power source device control. For example, when changing from a wired to a wireless network, the boot-time of APs must be considered (Figure 3, #10). The Network Management Layer waits for the AP to finish booting before it powers down the hub to the network resource where it was previously connected. This powering down of the hub forces the network resource to connect to the AP. The time when two network interfaces are available (after booting of the AP but before powering down the hub) allows a handoff between the two connections to occur.

3. **Run network scenarios**. The Control Layer is responsible for running network scenarios. It combines multiple actions of the Network Management Layer and the Verification and Validation Layer.
    - *Randomly produce and reproduce network scenarios.* The Control Layer or the user can specify a seed to generate a pseudo-random series of network scenarios. The seed from an earlier scenario can be used to recreate that series. The log from a series logs a seed at the start of each network scenario. Any of the logged seeds can be used to reproduce the sequence from that point forward in the log.
    - *Obtain results from multiple layers and combine into one report.* From the Single Control Point (Figure 3, #11), the Control Layer drives all other framework layers and obtains their results to include in the final report.
    - *Control all devices from one location.* The Control Layer communicates with the Validation and Verification Layer and the Network Management Layer (Figure 3, #12) via the multiple interfaces in Single Control Point.
    - *Coverage Analysis.* The Single Control Point reports coverage achieved at scenario completion.

Example Network Scenario Sequence

The following network scenario sequence illustrates one practical instance of the DND framework. To run a network scenario that transitions a mobile client from a wired network to a wireless network on a different subnet, the user would choose from a list of scenarios in the Control Layer (Figure 3, #11). Assuming this network must be created, the Control Layer resolves the particular case into a simple human-readable network request.

Based on the request, the Network Management Layer (Figure 3, #6) appropriately organizes the Physical Network Layer components into the configuration. The physical configuration contains mobile node laptops connected to a wired local area network (LAN). The Validation and Verification Layer confirms that the mobile nodes are reachable.

The Control Layer analyzes the results from the Network Management Layer to determine if the transition (wired subnet to wireless subnet) is possible. The Control Layer sends commands to the Network Management Layer to transition from wired to wireless. The first command powers on an access point (Figure 3, #10). After the access point has initialized, the hub that joins the mobile nodes to the VLAN is powered down (Figure 3, #9). The Validation and Verification Layer runs validation tests to verify that the mobile nodes are reachable on the wireless subnet. The Control Layer (Figure 3, #11) determines the success or failure of the network scenario and completes the sequence.

Practical Implications

The DND framework implementation discovered memory leak, continuous operation, firmware, roaming, and IP address defects that would have been hard or impossible to discover using traditional manual testing methods. Additionally, the DND framework implementation allowed reproducible test scenarios to aid debugging efforts in several cases, especially where manual testing was unable to reproduce similar scenarios.

The DND framework implementation also increased test coverage. For example, the DND framework implementation ran for 160 hours of continuous operation testing (non-stop) with 14,869 transitions occurring over hundreds of network configurations (approximately five network configurations every 70 minutes) across eight mobile nodes. This averages approximately 1858 transitions per node which approximates over a year (371 days) of usage by a user who averages five transitions a day. To accomplish similar test coverage in a manual environment and assuming two network configurations a day with 60 transitions (an extremely aggressive pace considering the time to manually set up and debug the network, verify results across all clients and collate data), this kind of coverage would require at least 248 days.

BENEFITS AND USES

The flexibility of the architecture allows a wide variety of uses by users with a wide variety of skill levels. IT technicians could use it to deploy new networks or network segments. The Verification and Validation Layer of the architecture could also provide valuable network status information such as network traffic patterns, downed network segments, network saturation, and network health. Marketing personnel could use it for presentations to quickly set up networks and move through different scenarios. Customers out in the field could provide product support engineers with the specifics of their network environment, and the product support engineers could then use the framework to quickly recreate the customer environment to resolve any issues. Validation engineers could use this architecture to automate test cases for increased testing coverage and log the results of tests on a wide range of real (non-simulated) networks. Validation and development engineers could use the framework to reproduce specific scenarios for debugging work. In each of these cases, no specific network component knowledge is needed to quickly and accurately configure a real-world network according to exact requirements. In addition, the cost in terms of resources, time, the need for network experts, and money is dramatically reduced by the adoption of the Dynamic Networks on Demand (DND) framework.

CHALLENGES AND NEXT STEPS

The next steps in developing the Dynamic Networks on Demand (DND) framework are to identify and extend the framework to handle the wider variety of scenarios required by the evolving needs of wireless networking projects.

For instance, creating interference of multiple signal strengths by access points (APs) would be useful for testing power optimizations and handoffs at boundaries as well as for testing behavior as signal strengths increase/decrease. The proposed solution would integrate an attenuator into the framework to programmatically initiate AP handoffs and increase/decrease signal strengths without actually moving components. The attenuator could also be used to address the challenge of introducing interference onto the wireless network.

Another example, stress testing of the AP with a maximum number of clients, is useful for test scenarios that examine behaviors of networks and clients in over-utilized environments. Solutions are being explored to load the network with a maximum number of clients without actual hardware, by using a simulator.

CONCLUSION

The framework is compelling for a variety of reasons. From a technical standpoint, the layered framework is extensible and reusable, and minimum effort is required to add new resources. Its modular design is easy to understand and enables easy addition of features. It greatly reduces manual operations (plugging and unplugging wires or devices, configuring devices), which saves network configuration turnaround and downtime, and it also allows a maximum number of network configurations to be achieved with a minimum hardware set.

From a non-technical viewpoint, by abstracting the network concepts and configurations, the architecture allows network operations to become more accessible as networks become deployable, without requiring technical knowledge about all network components.

From either view, the architecture provides support for numerous uses that require accuracy in network tasks while at the same time saving time, money, and resources.

ACKNOWLEDGMENTS

The team acknowledges Allen Sampson for nurturing a quality group that fosters test innovation and automation, Denise Shoup for her work on the Control Layer, Mike Kasun for his test setup and execution, and the support of Mike Andrews, Ranjit Narjala, and Prakash Iyer.

REFERENCES

[1] Robert V. Binder, *Testing Object-Oriented Systems Models, patterns, and tools,* Addison Wesley Longman Inc., Reading, Massachusetts, (c) 2000.

[2] Cem Kaner, Jack Falk, Hung Q. Nguyen, *Testing Computer Software, 2nd Edition*, John Wiley & Sons, (c) 1999.

[3] Boris Beizer, *Black-Box Testing: Techniques for Functional Testing of Software and Systems,* John Wiley & Sons, (c) 1995.

AUTHORS' BIOGRAPHIES

Amber Sistla is a network software quality engineer in the Network Architecture Labs in the Corporate Technology Group at Intel Corporation. Her technical interests include networking, wireless technologies, and test automation. She received her B.S. degree in Computer Science from Brigham Young University. Her e-mail is amber.sistla@intel.com

Jeremy Rover is a network software quality engineer in the Network Architecture Labs in the Corporate Technology Group at Intel Corporation. His current interests include wireless networking, dynamic network reconfiguration, and test automation. He received his B.A. degree in Computer Science from Linfield College. His e-mail is jeremy.rover@intel.com

Asha Keddy is a software quality manager in the Network Architecture Labs in the Corporate Technology Group at Intel Corporation. Her technical interests include network technologies, object-oriented design, software architectures, and innovative test automation. She obtained her B.E. degree in Computer Engineering from Bombay University India and her M.S. degree in Computer Science from Clemson University. Her e-mail is asha.r.keddy@intel.com

High-Throughput Wireless LAN Air Interface

Boyd Bangerter, Corporate Technology Group, Intel Corporation
Eric Jacobsen, Corporate Technology Group, Intel Corporation
Minnie Ho, Corporate Technology Group, Intel Corporation
Adrian Stephens, Corporate Technology Group, Intel Corporation
Alexander Maltsev, Intel Communications Group, Intel Corporation
Alexey Rubtsov, Intel Communications Group, Intel Corporation
Ali Sadri, Intel Communications Group, Intel Corporation

Index words: High-Throughput WLAN, 802.11n, OFDM, MIMO, Adaptive Bit Loading

ABSTRACT

This paper highlights the major candidate technologies for unlicensed band Wireless Local Area Network (WLAN) air interface performance improvements and discusses the advantages and disadvantages associated with each technology. We also explore the improvements in digital and radio frequency process technology that will contribute to the commercialization of high-throughput wireless LAN systems.

INTRODUCTION

Wireless Local Area Network (WLAN) technology is a fast-growing segment in the computing and communications equipment market, and the technology is evolving at a rapid pace. This evolution is happening as a result of coordinated standards efforts as well as innovation by leading-edge companies. It is taking place on many fronts: improvements in antenna systems, radio frequency (RF) components, modulation schemes, medium access control (MAC) mechanisms, and security mechanisms. Today, the air interface of standards-based WLAN equipment supports data rates as high as 54 Mbps. Designers of future WLAN equipment are attempting to specify an air interface that will achieve data rates in excess of 250 Mbps.

A variety of challenging tasks face designers of future air interfaces for WLAN systems. First, in order to serve a worldwide market, high-performance modes must be designed for more frequency bands than previous generations of WLAN equipment. Second, future designs must provide robust methods for co-existence with legacy equipment, and in many cases must also provide methods for backwards compatibility. Third, designers must select from a wide variety of performance-improving candidate technologies. The stand-alone performance characteristics of many candidate technologies are well understood; however, it is a challenge to understand how these technologies interact with each other as part of a complete system. Fourth, many of these candidate technologies require computationally intensive signal-processing algorithms and more complex radio frequency (RF) systems. Designers are faced with a large array of tradeoffs in order to find the optimum mix of performance, price, and power consumption.

THE CHALLENGES OF HIGH-SPEED DESIGN IN UNLICENSED BAND WIRELESS CHANNELS

Wireless channels present some challenging engineering problems that are not found in wireline systems. Most Wireless Local Area Network (WLAN) systems use omnidirectional antennas that provide good coverage but don’t concentrate the transmitted power to the intended user. This also means that because signal energy is scattered and reflected from objects in the environment, components of the signal arriving at the receiver are spread out over a longer period of time than is desirable. Since the frequency spectrum used by WLAN devices is unlicensed, other devices may be attempting to use the same channel resources and thus create interference. The challenge is then to provide a high-performance, reliable data link that can operate with restricted receiver power levels, severe channel fading due to multipath reflections, and interfering energy from other devices.

As with many technologies, it is useful to compare achieved performance with theoretical limits. In this case the theoretical limit is Shannon’s capacity, which is usually expressed as $C = B\log_2\left(1 + SNR\right)$

where B is the occupied bandwidth and SNR is the signal-to-noise ratio at the receiver. C is then the theoretical maximum rate in bits-per-second that can be achieved in that channel with no transmission reception errors.

This calculation has some important limitations, however. It is developed only for a single use of a memoryless channel impaired only by Additive White Gaussian Noise (AWGN). The result does not hold for receivers experiencing multipath fading or interference from other devices, which are both impairments that are often greater than the noise contribution in many wireless systems.

HIGH-THROUGHPUT WIRELESS LAN CANDIDATE TECHNOLOGIES

Fortunately, there are technologies that can take advantage of multipath in wireless channels. One such set of technologies involves the use of multiple antennas for both transmitting and receiving. Another such set of technologies involves wideband adaptive Orthogonal Frequency Division Multiplexing (OFDM). We describe both sets of technologies in the following sections.

MIMO Systems for High Throughput

Systems employing multiple antennas for both transmitting and receiving are often called multiple-input multiple-output (MIMO) systems. The primary advantage of MIMO systems over single-input single-output (SISO) systems is that in the presence of a “rich” multipath, a MIMO system with M antennas at the transmitter and M antennas at the receiver provides M times the peak throughput of an SISO system without increasing the frequency bandwidth. This is performed by dividing the channel into multiple “spatial channels” through which independent data streams can be transmitted. This technique is known as “spatial multiplexing.”

Additionally, a MIMO system can be used to increase diversity gain, where the slope of the Bit Error Rate (BER) curve when plotted against the Signal-to-Noise Ratio (SNR) can be dramatically increased, up to a maximum of MxMxL over a SISO system, where L is the number of effective, independent multipath components, or taps. This increase in diversity gain translates to an increase in range at a given BER or an increase in data rate at a given range up to the maximum rate of the original SISO link. Other possible gains, such as array gain or coding gain, are not discussed here.

Since our focus is on MIMO systems for high throughput, we first consider spatial multiplexing techniques and then consider diversity techniques to improve performance. This is analogous to a wideband, adaptive OFDM system, where the wide bandwidths are used to increase the peak throughput of the original link, and adaptive bit loading (a diversity technique) is added to improve performance. We categorize diversity techniques for MIMO systems into the following two categories. In the first category, the transmitter has no knowledge of the channel and uses a coding technique to achieve diversity over an ensemble average of channel realizations. In the second category, the transmitter has partial or full knowledge of the channel and uses this knowledge to increase diversity gain. In both categories, we presume the receiver has knowledge of the channel, derived from training sequences, in order to separate data from the multiple spatial channels.

The first category is conceptually the simplest. A transmitter simply encodes the bits over space and over frequency and transmits these bits over multiple spatial channels. The receiver then separates the symbols from the multiple spatial channels and decodes the bits. An example of an architecture belonging to the first category is depicted in Figure 1.

The second category of diversity techniques is more complex because in order for a MIMO transmitter to gain full or partial knowledge of the channel, one of the following two possibilities must occur: (1) knowledge of the channel at the receiver is simply turned around and used at the transmitter, and no feedback is necessary or (2) knowledge of the channel at the receiver is fed back to the transmitter using a handshaking protocol. An example of (1) is depicted in Figure 2, and an example of (2) includes adaptive bit-loading techniques, described in later sections of this paper.

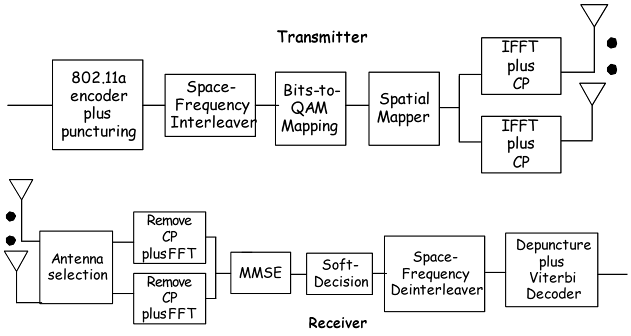

Figure 1: MIMO system without knowledge of the channel at the transmitter

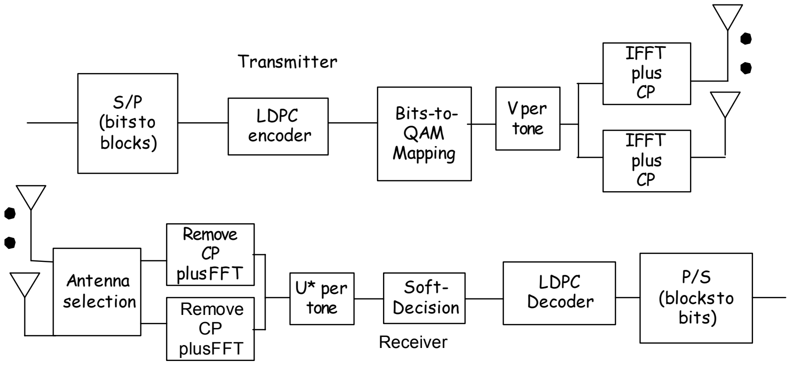

Figure 2: MIMO system with knowledge of the channel at the transmitter, but without feedback

MIMO Simulation Results

Figure 3 depicts the mean capacity vs. SNR (Shannon limits) of MIMO systems with different numbers of transmit and receive antenna elements. This simulation uses a random Rayleigh channel model. From this figure, we see that the capacity of the MIMO system with equal numbers of transmit and receive antenna elements increases almost *directly in proportion to the number of transceiver antenna elements* in comparison with usual SISO communication systems *without increasing the total transmitted power or frequency bandwidth*.

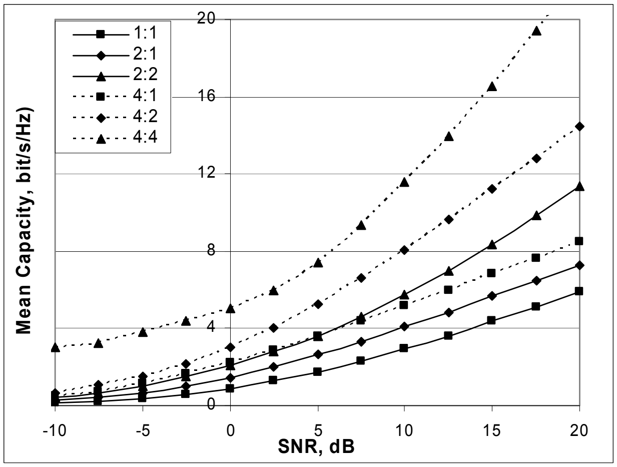

Figure 3: Mean capacity vs. SNR for MIMO systems

In Figure 4, the system depicted in Figure 1 was simulated with M_t transmit antennas and M_r receive antennas. Of these M_r receive antennas, M_c of the "best" antennas were chosen, where a Shannon capacity metric was used to determine the best antennas. This is denoted in the legend using the following notation: (M_r) M_c x M_t. We can see from this plot that we can achieve twice the throughput of a switched-diversity SISO (2) 1x1 system with a (4) 2x2 system, with a loss in SNR of 1 dB.

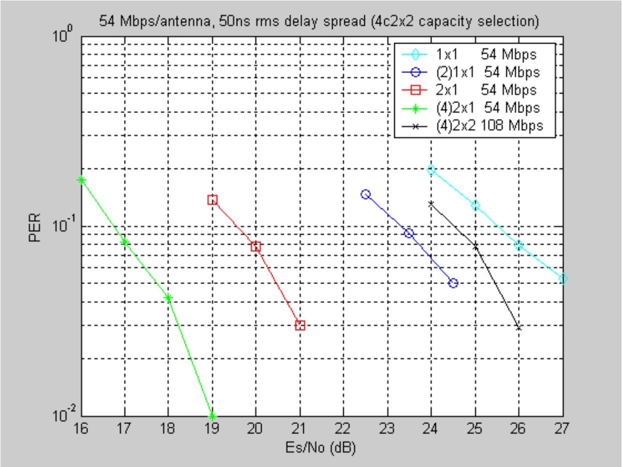

Figure 4: Simulation results for MIMO system depicted in Figure 1

In Figure 5, two comparisons are shown. The first comparison is between the techniques with and without channel knowledge at the transmitter: open loop minimum mean squared error (MMSE) scheme and closed loop singular value decomposition (SVD) scheme. The second comparison is between two forward error correction codes: (1) the 802.11a convolutional code and (2) a low-density parity check (LDPC) code described in a later section. Receive diversity selection is applied in all curves, and the number of antennas available to select are specified in the parentheses in the legend. In the SVD scheme, the channel matrix H is decomposed into the multiple spatial channels using a Singular Value Decomposition (H=UDV*). An LDPC code is applied to each of the spatial channels. We can see that the "SVD+LDPC+SVD" packet-error-rate (PER) curves gain about 3.56 dB in SNR over the "MMSE+CC" PER curve at PER 10^{-1}. The curves also demonstrate that the substitution of convolutional code with LDPC code by itself delivers 2 and 1.5 dB for open and closed loops, respectively. We also observe that we can achieve twice the throughput of a switched-diversity SISO (802.11a) system at a gain of 1 dB in SNR. It should be noted that the LDPC block size is of 4 µs duration, which is similar in latency to the convolutional code. In the figure, the upper solid lines are for PER and the lower dash lines are for BER performance.

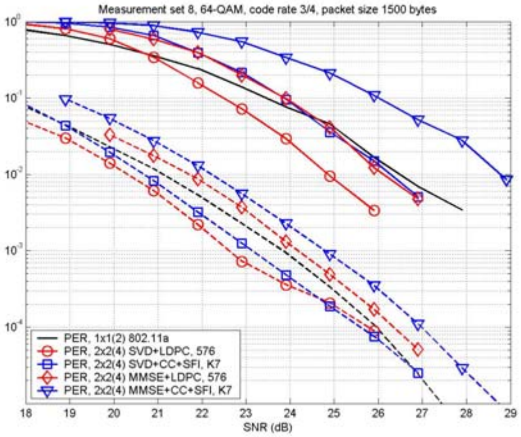

Figure 5: Simulation results for the MIMO system depicted in Figure 2

Wideband Adaptive OFDM

Another technology that may extend the capabilities of WLAN systems is adaptive OFDM. OFDM is the modulation currently used by 802.11a and 802.11g systems. OFDM works by dividing up a wideband channel into a larger number of sub-channels. By placing a subcarrier in each sub-channel, each subcarrier may be modulated separately depending on the SNR characteristics in that particular narrow portion of the band. As the channel varies over time, further adaptations can be made on each subcarrier in order to continually optimize the data-carrying capacity of the channel.

One possible approach to achieving higher WLAN data rates is to apply adaptive bit and power-loading techniques while simultaneously increasing the frequency bandwidth of signals from 20 MHz (as in 802.11a) to multiples of 20 MHz up to 80 MHz. Such an approach will result in capacity gains from both the adaptive techniques within the band as well as the raw linear increase in data rate expected with the increase in bandwidth.

Fast Link Adaptation With Adaptive Bit and Power Loading

Much research has gone into achieving the potential channel capacity using adaptive modulation, subcarrier power allocation, and coding techniques for OFDM systems. The underlying idea behind these methods is to exploit "good" and "bad" subcarriers in such a way as to maximize the data-carrying ability of the channel. This approach is often referred to as the "water-filling" approach. In the water-filling approach, more power and higher order modulation can be put onto subcarriers with larger SNRs, and lower SNR subcarriers will receive less power and lower order modulation up to a certain threshold, after which the subcarrier should simply be turned off, or "punctured."

A few different types of adaptive bit and power-loading algorithms have been proposed and analyzed. Some algorithms, such as those proposed by Chow *et al* [1] and Fischer with Huber [2], are suboptimal in the sense of the maximization of the SNR margin for a target data rate. Other versions of a finite granularity loading algorithm were proposed by Cioffi *et al.* [3] for transmission over asymmetric digital subscriber lines (ADSL). This algorithm maximizes the data rate at a given SNR margin that guarantees the BER is less than a certain target level. This algorithm turns off some of the subcarriers and divides the remaining (active) subcarriers into fixed subsets in accordance with either water-filling or uniform power loading. In each subset, subcarriers are given the same combination of "modulation+encoding," and then rescaling of subcarrier powers is performed. Hanzo *et al.* [4] made a significant contribution to refining adaptive bit and power-loading algorithms for duplex OFDM wireless links.

We propose and investigate different types of pragmatic adaptive bit and power-loading schemes that guarantee a target PER for wireless OFDM systems. They are suitable for duplex time division multiplexing (TDM) communication between two wireless stations, when they are acting over a slowly varying frequency-selective channel. Suppose that a first OFDM station (STA1) transmits data to a second OFDM station (STA2). During an initial handshake period (RTS/CTS exchange), STA2 measures channel characteristics observed when receiving the RTS frame and calculates an appropriate bit and power-loading allocation. STA2 transfers the results of this calculation to STA1 in the response (CTS) packet. STA1 will subsequently modulate subcarriers in accordance with the received modulation parameters. STA2, knowing these parameters in advance, will be able to receive and process the OFDM symbols. This closed-loop adaptation mechanism provides fast link adaptation that can follow time-varying channel characteristics.

We have developed a number of pragmatic finite-granularity adaptive bit and power-loading schemes that differ in throughput performance, complexity, and the amount of feedback service information. Three schemes are described below.

The Adaptive Bit and Power-Loading (ABPL) Scheme

This scheme exploits the water-filling principle for initial power allocation and determination of "no transmission" (non-active) subcarriers. The number of subcarriers in each subset with a uniform modulation plus encoding type is recursively determined. Then SNR pre-equalization is performed within each subset in order to guarantee a target fixed BER for each subcarrier. During this procedure, surplus power removed from higher-order modulation subsets is added to lower-order modulation subsets and is also used for turning on subcarriers that were initially turned off. At all stages this scheme is constrained by FCC requirements on peak power spectral density (PSD). It is the most computationally complicated scheme and has the best throughput performance. At the same time, however, it requires the largest amount of feedback service information, and it is very sensitive to channel state estimation errors.

The Adaptive Bit-Loading (ABL) Scheme

The ABL uses the same transmit power for all active subcarriers. It guarantees a BER less than some target threshold for each subcarrier and uses a "no transmission" mode for very bad subcarriers. The ABL scheme has some throughput degradation compared to the ABPL scheme, but it needs less feedback service information. Simulations have shown that the ABL scheme is more robust to impairments such as inaccurate channel state estimation and to interference and time variation of the wireless channel caused by Doppler spread.

To decrease the amount of feedback required, we investigated modifications of these schemes that allocate subcarriers the same parameters in groups of size 2, 4, or 6. In each group, subcarriers have the same power and modulation+coding type that are determined on the basis of effective noise power per subcarrier in the group.

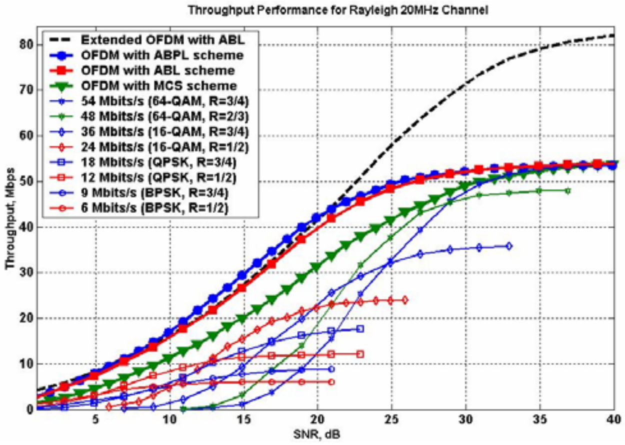

Figure 6: Simulation results for a static random Rayleigh 20 MHz channel model

Simulations show that a 2-subcarrier grouping has negligible throughput performance degradation (less than 0.3 dB in equivalent SNR); a 4-subcarrier grouping leads to performance degradation of about 1 – 1.5 dB, and a 6-subcarrier grouping gives substantial degradation in throughput performance (about 2 – 3 dB). A 2-subcarrier grouping halves the amount of feedback information required and is recommended for a high-throughput WLAN operating in a typical indoor environment.

Optimized Modulation and Coding Selection (MCS)

Another fast adaptation scheme is when all subcarriers are given the same modulation and power, based on channel state feedback. The optimal selection of modulation and coding type is made on the basis of the "momentary link performance" i.e., short-term SNR or PER estimates, and also a priori knowledge of momentary PER vs. momentary SNR performance curves [5]. We have investigated one algorithm for a modulation and coding selection (MCS) scheme. It shows a throughput degradation compared to ABPL or ABL. However, it requires a minimum amount of channel state feedback and gives a 1.5 – 2.5 dB throughput gain over a conventional long-term MCS strategy (see Figure 6).

Extended Bandwidth OFDM Simulation Results and Backwards Compatibility

Throughput and PER performances for the adaptive loading schemes were investigated for a 20 MHz OFDM system for different channel models, including models with Doppler spread. Results of the simulations for the static random Rayleigh channel model [6] are shown in Figure 2. Four modulations (BPSK, QPSK, 16-QAM, and 64-QAM) and three convolutional code rates (802.11a convolutional code with R=1/2, 2/3, 3/4) were considered. A 3-bit soft decision Viterbi decoder was built into the OFDM system model. The length of the data frame was 1000 bytes, and the target PER was 1%. Ideal carrier frequency estimation, timing synchronization, and channel state estimation were assumed.

It can be seen from Figure 6 that the ABPL scheme has about a 6 – 7 dB throughput improvement over standard OFDM. The ABL scheme gives about a 5 – 6 dB throughput gain over standard OFDM. The MCS scheme has a 1.5 – 2.5 dB throughput gain compared to a standard OFDM.

Extending Modulations, Coding Rates, and Frequency Bandwidth for HT WLAN

It is clear from Figure 2 that for a frequency-selective Rayleigh channel using any adaptive bit and power-loading scheme, a significant performance increase over conventional modulation is achieved, but there is no increase in the maximum achievable data rate (54 Mbps). We now consider two ways to enhance the maximum data rate in the context of single-antenna systems.

The first way is to use higher order modulations and/or higher coding rates to improve spectral efficiency. We have investigated adaptive loading schemes in conjunction with an extended set of modulation types (including 256-QAM) and coding rates (R = 7/8). This can achieve a data rate of up to 84 Mbps within a 20 MHz channel. The throughput results for the ABL scheme using the extended modulation and code-rate set is shown in Figure 6 by the black dashed curve.

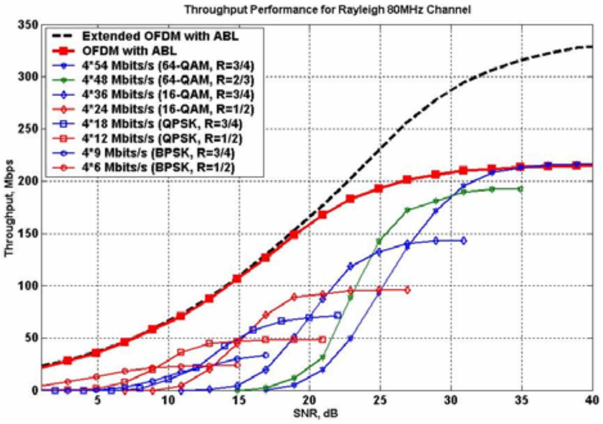

Figure 7: Simulation results for static random Rayleigh 80 MHz channel model

The second way is to use increased channel widths. We investigated using up to four 20 MHz channels in one U-NII band in conjunction with the adaptive loading schemes. For example, simulation results for an OFDM system with ABL, operating in a static random Rayleigh 80 MHz channel model [5], are shown in Figure 7. It can be seen that the ABL scheme gives about a 6 dB throughput gain (see bold red curve with squares) over conventional OFDM. In an extended bandwidth mode, the OFDM system throughput can achieve 300 Mbps (see black dashed curve).

Throughput Versus Range Performance for Extended Bandwidth OFDM Systems With an ABL Scheme

Extended bandwidth OFDM leads to decreased transmit power spectral density because of FCC limits on the total transmit power of wireless devices. Increasing frequency bandwidth by a factor of four results in a decreased SNR per subcarrier by the same factor. This causes range performance degradation in extended bandwidth OFDM systems.

At the same time, using these adaptive bit and power-loading schemes in extended bandwidth systems leads to economic use of available transmit power because only good subcarriers are allocated power, and bad subcarriers can be turned off.

We investigated throughput versus range for extended bandwidth OFDM systems using adaptive loading schemes. Figure 8 shows simulation results for an indoor environment with an exponential channel path loss model and one additional wall (exponent equal to 2.15) [6]. We used the ABL scheme in the extended mode (256-QAM and R = 7/8) for OFDM systems with various frequency bandwidths (20, 40, 60, 80 MHz) and different total transmit power. OFDM systems with ABL schemes demonstrate both maximum throughput improvement and range extension.

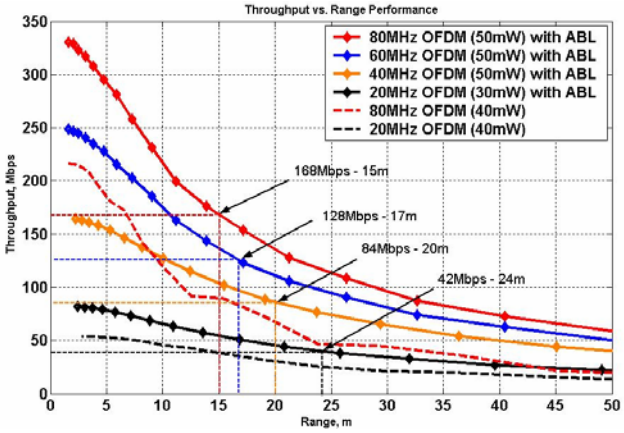

Figure 8: Throughput versus range performance for extended bandwidth OFDM systems with ABL scheme

Degradation of ABL Throughput Performance Due to Time Delay of Data Transmission

To study ABL throughput performance degradation in a non-stationary environment we used an indoor channel model with Doppler spread. The impulse response of a frequency-selective channel with Doppler spread was modeled as a tapped delay line where tap coefficients were independent Rayleigh fading quantities (Jakes model) [7]. In this channel model, the tap coefficients vary slowly because of mobile and scatterer motion. To illustrate this, Figure 9 shows variations of the channel transfer function for one random 20 MHz channel realization during a 10 ms time interval.

A variation in the channel transfer function during packet exchange may lead to performance degradation for any adaptive bit and power-loading scheme.

This channel model was used to investigate Doppler spread affects on the performance of an OFDM system with ABL. The channel is measured by STA2 receiving (for example) an RTS frame from STA1, and the optimal bit and power loading is calculated at the end of the RTS frame. These results are sent to STA1 in some response frame (for example, CTS) and applied by STA1 during transmission of the subsequent data packet to STA2. This data packet from STA1 is sent through a channel that is changed due to Doppler spread and exchange time delay.

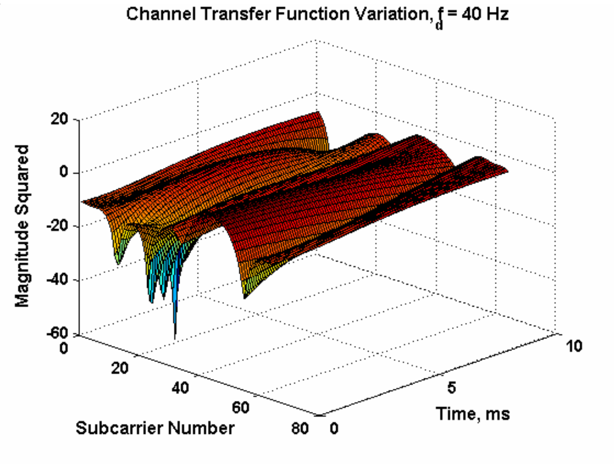

Figure 9: Channel transfer function variation due to Doppler spread with f_{dmax}= 40 Hz

We have investigated the throughput performance degradation of an OFDM system with ABL versus time delay between channel state measurement (by STA2) and ABL data packet transmission (by STA1). The length of the data packet was 1000 bytes.

Simulations of the OFDM system with ABL have shown that even for the modeled severe Doppler spread (f_{dmax} = 40 Hz), throughput degradation no worse than 1 dB is seen when time delay is about 2 – 2.5 ms.

Channel Feedback Information on Backwards Compatibility Issues

Increasing the frequency bandwidth of OFDM signals and exploiting adaptive bit and power-loading techniques requires development of a packet structure for the extended bandwidth OFDM physical layer (PHY). A new packet structure is required that can overcome the frequency selectivity and non-stationary behavior of the wider channel, and that can provide backwards compatibility with existing IEEE 802.11a systems operating in the same frequency band.

We propose to respect the existing time- and frequency-related parameters associated with the OFDM PHY layer convergence protocol (PLCP) of IEEE 802.11a. For example, all simulations for extended bandwidth OFDM systems were done assuming that OFDM data symbols are generated by replication in the frequency domain of conventional 20 MHz OFDM symbols.

To reduce interference from legacy IEEE 802.11a devices, we suggest the use of a “compatibility preamble” mechanism. The compatibility preamble may be constructed from one to four IEEE 802.11a PLCP headers (including preamble and signal field) spread in the frequency domain. This provides support for the coexistence of extended bandwidth OFDM systems with legacy wireless devices. The compatibility preamble allows legacy devices to receive the beginning of an extended bandwidth OFDM transmitted packet. Although they cannot decode the remainder of the packet, they honor the duration specified in the legacy PLCP header, thereby providing protection to the extended bandwidth OFDM packet from legacy interference. The compatibility preamble can also be used for the conventional receiver tasks of channel estimation, timing synchronization, and frequency offset estimation.

In order to exchange feedback service information between extended bandwidth OFDM stations needed for ABL, we propose to insert into the OFDM packet an additional OFDM PLCP header that can be understood only by next-generation WLAN devices. It contains channel state feedback and is modulated with the most robust modulation and encoded with the slowest code. 3 to 5 OFDM symbols of feedback (BPSK with R=1/2) should be enough for successful operation of extended bandwidth OFDM systems with ABL. The OFDM PLCP headers also contain any additional parameters that are required to manage the packet exchange.

ADVANCED CODING

Nearly all wireless systems employ Forward Error Correction (FEC) techniques to correct the numerous transmission errors that occur in wireless channels. As shown previously, Shannon provided a means to calculate a theoretical maximum transmission rate for Additive White Gaussian Noise (AWGN) channels at which error-free transmission could be achieved. It took almost fifty years from the development of Shannon’s theoretical limits until FEC codes were developed that could reasonably approach those performance levels. Only recently are these types of codes making it into the commercial market, and research continues into developing more efficient solutions.

The existing IEEE 802.11 wireless standards all specify a well-known convolutional FEC code, and the Viterbi algorithm is universally used for decoding these codes. At the time of the development of the current standards this was the most practical solution considering cost, complexity, power consumption, and decoding latency. Unfortunately convolutional codes and the Viterbi algorithm leave a significant amount of power efficiency unclaimed with respect to the theoretical capacity limits.

About ten years ago the FEC community experienced a revolution with the discovery of Turbo codes, which were the first practical codes developed that could come reasonably close to theoretical capacity. Even after ten years of additional research, Turbo codes have seen limited acceptance in certain applications due to their complexity and decoding latency. One of the difficulties with Turbo codes is a loss of performance when the size of the code block is reduced. For streaming applications like broadcast video, this is not a problem, but WLAN packet sizes are often very small and transmissions as simple as packet acknowledgements require link reliabilities comparable to long packets. This degradation of performance for short blocks as well as their implementation complexity makes it difficult to successfully apply Turbo codes to WLAN applications.

The discovery of Turbo codes, which use iterative decoding to achieve much of their performance, renewed research interest in another family of FEC codes known as Low Density Parity Check (LDPC) codes. Although this type of code was initially investigated in 1963 by Gallagher [8] little research was devoted to them until the publication of Turbo codes sparked interest in iterative techniques. Recent research has shown that carefully designed LDPC codes do not suffer the same performance degradation for short data blocks as Turbo codes. The underlying codes in LDPC codes are very simple parity-check relationships, so the potential for complexity reduction compared to Turbo codes is significant.

Figure 10 shows a performance comparison of a type of LDPC code developed by Intel for possible use in WLAN systems against the Viterbi decoder currently used in 802.11 systems. In this case, blocks of 2193 data bits are used to generate 731 parity bits that are appended to make a 2924-bit codeword (i.e., a (2924, 2193) code). This provides a code of the same rate as the R = ¾ convolutional code used for comparison. Longer data packets would be coded by concatenating successive code words. Shorter packets would be accommodated by using code shortening, where the encoder and decoder assume the remainder of the data block is zeros. Since the assumed zero-padding values are known, the unused portion of the shortened code word does not need to be transmitted.

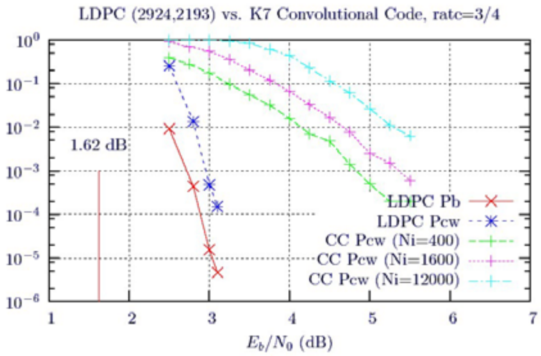

Figure 10: Comparison of Bit Error Probabilities (Pb) and Code Word Error Probabilities (Pcw) for the K = 7 convolutional code used in the current 802.11 standards and a proposed (2924, 2193) LDPC code

Shortening the LDPC code for smaller blocks maintains a significant improvement over the current system. Capacity for an R = ¾ code is shown at Eb/No = 1.62 dB. It can be seen that the LDPC code is operating less than 1.5 dB from capacity. Performance for the convolutional code is shown for data blocks of 400, 1600, and 12000 bits.

The performance increase shown in Figure 10, where the LDPC codes require about 3 dB less transmit power to achieve Pcw = 10^{-5}, can be used to increase range or to increase the selected modulation order for higher throughput.

HIGH-THROUGHPUT MAC CONSIDERATIONS

Why do we need to modify the MAC at all? Isn't getting high throughput largely a PHY issue? The answer is "no" because while the PHY data rates are increasing significantly, the PHY overhead is not decreasing by the same factor. Therefore, throughput becomes increasingly dominated by overhead (radio turnaround times, modem pipeline delay, preamble, PLCP headers, etc.).

In order to make the most efficient use of the wireless medium, we need to take a system view and design the operation of the MAC so that it effectively uses and manages the services of the PHY while providing an unmodified interface to the higher layers.

Interface to Higher Layers

The MAC provides a connectionless (so-called UNITDATA) interface to higher layers. This means that each packet sent down from the higher layers is treated by the MAC independently of all others. The MAC preserves order between any source/destination address pair.

While there is an increasing amount of knowledge in the higher layers used to manage the MAC through its management interfaces, from the data point of view, it is still a "wireless Ethernet."

MAC Improvements

The improved MAC has two main jobs to do: (1) it manages the new features of the PHY and (2) it improves throughput by aggregation and scheduling.

In conventional 802.11 systems, the MAC can manage the PHY by adapting its transmit rate to observed conditions (you might see this as an "automatic" rate setting on your control panel). There are a number of things it can attempt to observe to make these decisions, including packet error rate. However, without explicit communication from its peer about receive conditions, its decisions are based on possibly flawed assumptions. Likewise (but less frequently attempted), conventional systems may try to automatically adapt fragment size.

The WLAN system described here gains significant throughput improvements by adapting its transmit parameters to closely match the properties of the link between two devices. The more capacity the link has, the more throughput the system will obtain. It does this through a training process that can be repeated as often as necessary to keep training data fresh.

The other new main job the MAC has is to create bursts of packets addressed to the same destination. It sends these using the 802.11e Block Ack mechanism thereby significantly reducing overhead.

MAC Tradeoffs

In order to create bursts, the MAC has to buffer traffic by destination address until some criterion (e.g., burst length or age) makes it eligible for transmission. This creates a scheduling problem. We have to find a balance between fairness, delay, and throughput. This is an ongoing topic of research.

A next-generation MAC will also support 802.11e Quality of Service (QoS) enhancements (this is necessary in order to use 802.11e Block Ack). 802.11e supports negotiation of the QoS attributes of a traffic stream through a "traffic specification" (TSPEC). The TSPEC allows a next-generation WLAN device to respect traffic-specific delay limits in its formation of bursts.

Another tradeoff the MAC can manage is the liveness of training data. The MAC obtains training data through an exchange of training frames with a particular destination. This training exchange is an overhead. The MAC can track the age of training data and decide when it is likely to be out of date.

When the MAC decides to transmit to a destination, it can optimize throughput based on the age of the training data and the amount of data to be sent. For example, it is not worth the overhead of re-training just to send small amounts of data.

Coexistence Options

It is the primary job of the MAC to manage access to the wireless medium.

We want to make efficient use of the wireless medium, but we want to be fair to existing users.

In most deployment scenarios, next-generation access points (APs) will be added to create a new network or to overlay an existing network. In these cases, the best coexistence option is for the new APs to select an operating channel that avoids creating any overlap with legacy 802.11 devices sharing the same band. In the case of enterprise deployments, frequency planning would make this happen. In the home environment, the lower installed density makes it very likely the new APs can automatically find a clear operating channel.

If the next-generation device cannot find a channel clear of legacy devices it can attempt to negotiate with legacy devices (e.g., using 802.11h signaling) to get them to relocate.

If this fails, the network has to operate in a so-called "mixed" mode that requires the use of a protection mechanism. It also requires that such devices be capable of sending and receiving legacy frames so that they respect medium reservation of the legacy devices. The protection mechanism uses the transmission of legacy frames (or compatibility preambles, as described above) so that the legacy devices respect either MAC-level or PHY-level medium "reservations."

A next-generation AP controls whether its stations operate in mixed mode. An alternative to mixed mode operation is for the AP to manage access to the wireless medium by preventing groups of stations from contending for access by setting their Network Allocation Vector (NAV), for example using beacons. This AP avoids the overhead of the protection mechanism and can control the amount of time the medium is used by legacy or high-throughput devices according to its own local policy.

CONCLUSION

This paper has shown that there are techniques that can improve throughput and range for next-generation wireless LAN devices. MIMO-OFDM techniques, fast link adaptation, and advanced coding schemes are likely to be used in next-generation WLAN products. These techniques neatly dovetail together to provide a system including both MAC and PHY layers that increases channel capacity and makes the most of this capacity.

Intel is contributing to and committed to work towards the development of this standard using these kinds of techniques.

ACKNOWLEDGMENTS

The authors acknowledge Alexei Davydov, Andrey Pudeyev, Vadim Sergeyev, and Sergey Tiraspolsky, who provided the simulation results for WB OFDM with fast link adaptation techniques. They also acknowledge Bo Xia for providing the LDPC simulation results and Sumeet Sandhu and Qinghua Li for contributing to the MIMO simulations.

REFERENCES

[1]Chow, P.S., Cioffi, J.M., and Bingham, J.A.C., "A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels," *IEEE Transactions on communications*, V. 43, p. 773, 1995.

[2] Fisher, R.F.H. and Huber, J.B., "A new loading algorithm for discrete multitone transmission," *Proc. IEEE GLOBECOM'96*, p. 724.

[3] Leke, A. and Cioffi, J.M., "Transmit optimization for time-invariant wireless channel utilizing a discrete multitone approach," *Proc. of IEEE ICC'97*, V.2, pp. 954 – 958.

[4] Keller, T. and Hanzo, L., "Adaptive modulation techniques for duplex OFDM transmission," *IEEE Transactions on Vehicular Technology*, Vol. 49, Issue 5, Sept. 2000, pp. 1893 – 1906.

[5] Simoens, S. and Bartolome, D., "Optimum performance of link adaptation in HIPERLAN/2 networks," *Vehicular Technology Conference*, Vol. 2, May 2001, pp. 1129 – 1133.

[6] Rhodes, V.J. and Jacobsen, E.A., "Static propagation characteristics for 5 GHz Wireless LAN applications," (to be published).

[7] Jakes, W.C., *Microwave Mobile Communications*, John Wiley & Sons Inc., New York, 1974.

[8] R.G. Gallager, "Low Density Parity Check Codes," *IRE Trans. Information Theory*, Vol. IT-8, pp. 21-28, Jan. 1962.

AUTHORS' BIOGRAPHIES

Boyd Bangerter manages the System Architecture and Standards group in Intel's Corporate Technology Group. He joined Intel in 1993 and has managed various programs related to wired and wireless broadband access technologies for the past ten years. He holds an MBA degree from Brigham Young University and a BSEE degree from the California Institute of Technology. His e-mail is boyd.r.bangerter@intel.com.

Eric Jacobsen is a wireless systems engineer in Intel's Corporate Technology Group. He is Principal Investigator in the Advanced OFDM research group where his current focus is on low-cost high-speed WLAN solutions. He is a past member of the IEEE 802.16 Working Group and is an active member of the 802.11 High Throughput group.

Eric has 18 years of experience in research and development in signal processing and communications in radar, avionics, satellite, and network applications. He has BSEE and MSEE degrees from South Dakota School of Mines and Technology. His e-mail is eric.a.jacobsen@intel.com

Minnie Ho leads a group conducting Smart Antenna Research in Intel's Corporate Technology Group. She joined Intel in 2001 as part of the Ultra-Wideband Research group. Prior to joining Intel, she worked at four start-up companies on projects involving ADSL (technology standardized by ITU), broadband wireless access, fiber-optic, and ultra-wideband communications. She holds M.S. and Ph.D. degrees from Stanford University in Electrical Engineering, and a B.S. degree in Electrical Engineering from Princeton University. Her e-mail is minnie.ho@intel.com

Alexander Maltsev is a staff research scientist in Intel's Wireless Networking Group. He leads the WPD Advanced Development Russian team, which works on research and development of signal processing algorithms for wireless communication systems.

Alexander has over 30 years experience in the development of analog and digital signal-processing algorithms, synchronization systems, adaptive antenna arrays, adaptive active vibration and noise control systems; and space-time signal processing in non-stationary environments. He received the Candidate of Science (Ph.D.) and Doctor of Science degrees from Nizhny Novgorod State University (NNSU), Russia, in 1975 and 1990, respectively, all in Radiophysics. Since 1994 he holds the chair in Statistical Radiophysics at NNSU. His e-mail is Alexander.Maltsev@intel.com

Alexey Rubtsov is a researcher in Intel's Wireless Networking Group. He works in the WPD Advanced Development team where his current focus is on adaptive bit and power-loading algorithms for OFDM systems. He is a student member of the IEEE.

He received his B.S. and M.S. degrees with honors from Nizhny Novgorod State University, Russia in 1999 and 2001, respectively. He is currently pursuing a Ph.D. degree in Radiophysics at NNSU. His e-mail is Alexey.Rubtsov@intel.com

Ali Sadri received his Bachelors, Masters, and Ph.D. degrees in Electrical Engineering in 1989, 1991, and 2000, respectively from North Carolina State University. From 1990 through 2000 he was with IBM in the Research Triangle Park, NC, where he was initially a member of the voice band modem and multimedia team. During the last seven years of Dr. Sadri's work at IBM, he became the program manager for IBM telecommunications standards where he was leading and coordinating IBM contribution and strategy at the International Telecommunications Union (ITU) and Telecommunications Industry Association (TIA). Dr. Sadri later joined BOPS Inc. During the years of 2000-02 he was the director of Telecommunications Development where he led the effort to design a low-power programmable wireless LAN DSP core. In the past two years he has been an adjunct associate professor at Duke University and a lecturer at Oregon State University. He is currently working at Intel Corporation and leading the advanced WLAN development effort. His areas of research are adaptive modulation, Adaptive TPC, WCDMA and OFDM. His e-mail is ali.s.sadri@intel.com.

Adrian Stephens is a wireless MAC architect in Intel's Corporate Technology Group. His focus is on developing next-generation IEEE 802.11 standards. He is an active member of the 802.11 High Throughput group (currently chair of their usage model committee).

Adrian has over 20 years of product development and research experience working for industry and government. Adrian received a B.A. degree and a Ph.D. degree from Cambridge University. He is a member of the IEEE, the IEE, and is a chartered engineer. His e-mail is adrian.p.stephens@intel.com.

Seamless Connectivity to Wireless Local Area Networks

Allan Chin, Intel Communications Group, Intel Corporation
Ajay Gupta, Intel Communications Group, Intel Corporation
Ranjit Narjala, Corporate Technology Group, Intel Corporation
Venkata Vallabhu, Intel Communications Group, Intel Corporation

Index words: Wireless LAN, Adapter Switching, Automatic VPN Launch, Seamless Roaming

ABSTRACT

The Internet is clearly becoming more diverse in its types of constituent interconnected networks. Wireless Local Area Networks (WLANs) are becoming ubiquitous in today's technology-driven world. Quick and easy connectivity to these networks is a big step in achieving the objectives of continuous mobility, constant reachability, and consistent connectivity.

This paper describes the challenges associated with a client solution and discusses how the Intel® Wireless PROSet software could support quick and seamless connectivity to WLANs, including additional features that enhance the user experience. Some of these challenges and solutions are covered in this paper as follows:

1. Configuration and manageability of the WLAN adapter and the available networks, and Intel's Wireless Configuration feature.
2. Multi-homed systems with their associated challenges, and a proposed solution using the Adapter Switching feature.
3. Security implications of wireless networks, and the proposed automatic Virtual Private Network (VPN) launch feature.
4. Co-existence with other "smart clients" on the client machine.

We also provide a walk-through of a common usage scenario, and depict how this software could enable a mobile user to truly enjoy the advantages of being wireless. Finally, we highlight what lies ahead in the world of constant connectivity and continuous roaming.

INTRODUCTION

The Internet is clearly becoming more diverse in its constituent interconnected networks. Wireless 2.5G/3G licensed spectrum radio networks, unlicensed spectrum-based Wireless Local Area Networks (WLANs) and Personal Area Networks (PANs), and wired IP-based networks will soon converge and interconnect to provide a rich medium to drive voice-data convergence to IP. Internet-attached devices are becoming mobile and multi-modal (i.e., capable of connecting to more than one physical network type). Due to higher bandwidth availability and ubiquitous reachability, peer-to-peer communication will become more prevalent.

In this paper, we describe the challenges that are peculiar to WLANs, and go on to illustrate how the Intel Wireless PROSet solution could meet these challenges.

Unlike wired networks, one cannot physically connect a wire to a wireless network. The client needs to listen for signals coming from the access point (AP) in order to establish a connection. By automatically sensing 802.11 signals and intelligently auto-connecting the user to the "best" wireless network, based on a user-defined configuration in the form of profiles, the Wireless Configuration feature essentially eliminates the complexity associated with configuring and managing WLANs on the client.

Many laptop computers in use today contain both wired and wireless network interface cards, thus making them multi-homed. This means that users cannot deterministically select one interface to be used for all traffic when multiple interfaces are available. In this paper, we describe a solution that allows the user the ability to prioritize a list of media types; the software then intelligently and automatically selects, and presents for use, one particular adapter when multiple adapters may be present. We then touch on the possibility of using this platform approach to intelligently conserve battery power.

Intel's "Intelligent Roaming" concept defines two different roaming models:

1. Nomadic Computing: This is a usage model where a user easily connects to the best available network

without expectation of session continuity (for example, a persistent TCP connection or a persistent Virtual Private Network (VPN) session) across networks. Nomadic computing can be further broken down as follows:

Quick Connect: In this situation the user's mobile device (i.e., laptop) has only one physical network adapter available for use. As soon as the adapter (otherwise called the Network Interface Card, or NIC) becomes active (i.e., enters a link-up state), a connection to the target network will be automatically initiated if permitted by policy. A VPN session may be activated if necessary as determined by a profile associated with the physical NIC type and network identifier, such as the WLAN Service Set Identifier (SSID).

Switched Roaming: This is a refinement of the Quick Connect scenario and is applicable only when the mobile device has two or more physical network adapters that are available for use. In this situation, if connectivity is possible on more than one interface, the software will select and activate the interface that is preferred by the user as determined by a local policy.

2. Continuous Roaming: This is a further refinement of the connectivity types listed above, adding session continuity across IP subnets and across physical network adapters. This solution can also provide for maintaining session continuity across VPN sessions. Session continuity is obtained by allowing the mobile client to retain a persistent IP address while moving across different IP subnets, using the Mobile IP standard (RFC 3344).

This paper describes a solution that supports only nomadic computing. The long-term goal is to provide technology building blocks to enable heterogeneous roaming across global wired and wireless IP-based networks while accomplishing the objectives of continuous mobility, constant reachability, and consistent connectivity.

Due to the broadcast nature of the wireless medium, security is required at the link layer. However, Wired Equivalent Privacy, or WEP-based security for wireless LANs has proven to be weak; an accepted solution is to deploy security at the network layer using IPSec. This paper describes how secure layer-3 connections are established as part of connecting to a wireless network.

With WLANs being increasingly deployed in hotspots, there are several smart clients from various clearing-house vendors and hotspot operators. We describe how Wireless PROSet automatically detects other smart clients in the system and co-exists with them.

We provide the reader with a walk-through of a common usage scenario and depict how this software could enable a mobile user to truly enjoy the advantages of being wireless. We end by highlighting what lies ahead in the world of constant connectivity and seamless roaming.

Note: The architectural views expressed in this paper are solely that of the authors. The authors make no express claims or comment on either existing or future Intel products, including the PROSet framework, based on ideas conveyed on this paper.

The next four sections describe some of the challenges that are unique to the configuration and use of Wireless Local Area Networks (WLANs), from a client perspective.

CONFIGURATION AND MANAGEABILITY

The Problem

Wireless networks are fundamentally about mobility. Users move between a number of locations that offer connectivity to different networks by using varied security models, and sometimes by requiring different TCP/IP settings. To switch between these networks, a user is required to make various network configuration changes. A lot of these changes require an in-depth understanding of network configuration, TCP/IP settings, and security models. Some networks also require additional software, such as Virtual Private Network (VPN) clients, to be run in order to obtain complete connectivity.

Ad hoc networks, where clients connect as peers directly to each other, are also becoming more popular. They offer the benefits of being able to easily and quickly exchange data between peers without the need to set up any infrastructure. Setting up the client to act in this mode can be very challenging; switching between ad hoc mode and normal (or infrastructure) mode can also be intimidating.

Configuring and managing these wireless networks is thus a time-consuming, tedious, and error-prone process that only very sophisticated users are able to successfully accomplish. Since one cannot get around the need to make these changes, the need for a tool that makes these changes automatically is highly desirable, and in fact is necessary if a normal user is expected to be able to benefit from wireless connectivity.

Our Solution

Intel's Wireless PROSet software suite enables the user to configure and manage all supported Intel wireless (802.11a/b) network adapters present in the system.

Users can create "profiles" for the different networks they would like to connect to; each of these profiles contains all the information necessary to connect to a particular network. This suite also provides a wizard for quick and easy configuration of these profiles, and allows the user to prioritize the list of profiles. Another valuable feature in the software suite is the ability to import and auto-import profiles. This allows an IT administrator to create profiles for the company WLAN and then provide it to the employees for import, or push it onto the employee machines for auto-import.

The Intel Wireless Configuration service, a component included in this software suite, will use the information specified in these profiles to automatically and seamlessly connect to the best available network, including ad hoc networks. This saves users from having to determine what network they need to connect to, and it alleviates the need for users to constantly remember all the network configuration parameters and to change the settings every time they move to a different network.

The Details

Quite unlike wired 802.3 networks, a user who would like to connect to an 802.11 WLAN does not need to physically connect a network cable to use the network. Instead, a number of WLANs may be available and connected to by merely "listening" to beacons that are broadcast by WLAN access points (APs) and/or WLAN client peers operating in ad hoc mode. Connecting to one of these WLANs requires programatically providing the "name" (SSID) of the network that the client desires to be connected to, and also optionally providing the authentication and encryption parameters that may be required for the particular network.

The ability to connect to various WLANs without requiring any user intervention facilitates the concept of automatic and seamless network connection services.

Firstly, the Intel Wireless PROSet software enables the user to create a WLAN profile or set of profiles, each profile consisting of information necessary for automatically connecting to a network, i.e., a set of network identification parameters as well as authentication and encryption parameters. There are essentially two types of profiles, one for infrastructure mode and one for ad hoc mode.

When the wireless adapter is configured in infrastructure mode, it will connect to a wireless AP. In ad hoc mode, however, the wireless adapter will look for and connect to other wireless adapters that are within mutual communication range of each other. An ad hoc network is typically created spontaneously, and its most distinguishing feature is its limited temporal and spatial extent. This mode allows users to quickly and easily exchange data without requiring any infrastructure (such as APs) to be set up.

Once a set of profiles have been configured, the user may also optionally prioritize this list. The Intel Wireless Configuration service, which is a component of the Wireless PROSet software, will attempt to first connect to a network using settings in the user's "most preferred" profile. This is accomplished by populating a "scan-list" that comprises the currently available networks, and then running through the profile list to determine which one can be used. The process of mapping the user's WLAN profiles to available networks uses the information available from the scan-list.

The Adhoc Wizard, another component of the Wireless PROSet software, enables the use of the wireless adapter in ad hoc mode. With the Adhoc Wizard, a user can select, from a list of peers, a particular peer he wishes to communicate with. The Adhoc Wizard will set up communication with that peer. Once the necessary data transfer with that peer is complete, the user can close the Adhoc Wizard, which will result in a connection to an infrastructure network being re-established.

Intel's Wireless Configuration service goes beyond Microsoft's ZeroConfig by allowing the user to specify, as part of an infrastructure profile, TCP/IP settings and VPN parameters. The VPN settings specified in the profile allow a VPN tunnel to be automatically established, once network connectivity has been obtained, therefore obviating the need for the user to remember that a VPN tunnel needs to be established before using the wireless network.

Another difference between Intel's Wireless Configuration service and Microsoft's ZeroConfig service is the way in which "stealth" networks are handled. A stealth network is one where the network's primary identifier (its SSID) is not broadcast via 802.11 beacons. While Microsoft's ZeroConfig service gives a higher priority to non-stealth networks, Intel's Wireless Configuration service establishes priority solely based on the order of the user's profiles.

In addition, Intel's Wireless Configuration service provides these automatic and seamless connectivity features not only on Windows XP and Windows 2000,* but also on older versions of Microsoft's operating systems, such as Windows NT* and Windows 98,* thus allowing a variety of users the ability to enjoy the benefits of wireless networks.

Summary

Intel's Wireless PROSet software suite thus offers the following benefits:

1. It enables the user to create and prioritize a set of profiles for different WLANs. These profiles are used by the Wireless Configuration service to automatically connect to the best available network. Users therefore do not need to constantly configure the wireless adapter each time they move into a different network.
2. This software will work on almost all versions of the Windows operating system in use today, thus providing a variety of users with the ability to easily configure and manage their WLAN connections.
3. The Adhoc Wizard enables seamless connectivity to ad hoc networks. It automatically provides a list of available peers within the ad hoc network specified by the user, which helps users see the other peers they can connect to. It also connects back to an infrastructure network when the ad hoc connection is terminated, thus simplifying the entire connection process.

ADAPTER SWITCHING

Problems Associated with Multi-Homed Hosts

The standard for today's laptops is to come with at least two different types of network adapters: a wired adapter and a WLAN adapter. Most enterprises support both wired and WLAN connectivity. Users are also starting to install and use both wired and WLANs at home, along with their broadband Internet connections. As a result, the user may be simultaneously connected using both a wired and WLAN adapter, and is able to communicate using both these adapters.

To understand the problems that arise when the user is connected using more than one network interface, we need to take a closer look at how the Microsoft Windows operating system functions when multiple adapters are present. In general, Windows maintains a route-table for all installed network adapters (also known as interfaces), and this route-table is consulted to determine which interface an outbound packet should be sent out on. Whenever an adapter gets an IP address, Windows inserts a set of entries into this route-table. If more than one adapter is in a link-up state and acquires an IP address, each of these adapters will have a set of entries in the route-table that point to itself. When an application sends a packet out, the packet will first encounter the TCP/IP stack in the kernel. The TCP/IP stack will consult the route-table to determine which interface this packet should be sent out on. If there are multiple choices because of the multiple adapters that are available, TCP/IP will select one particular adapter, based on a set of internal algorithms, and the packet will be sent out over that interface. Subsequent packets from the same application that belong to the same stream will be sent out over that interface, as long as the interface remains in a link-up state and maintains a valid IP address. The user thus cannot deterministically ensure that all traffic will go out a particular "best" interface when more than one interface is available—the choice of the interface is up to the operating system.

Our Proposed Solution

Adapter Switching could step into the picture at this point, and allow the user the option to deterministically select one interface to be used for all traffic when multiple interfaces are available. The user can prioritize the adapters present in the system based on the adapter type. The Adapter Switching software will select the best adapter out of those available, depending on the user's preferences and present that to the operating system to use for all network connections.

The Details

As mentioned in the introduction, Quick Connect is a feature that will enable the user to quickly and easily connect to an available network when the user's mobile device has only one physical adapter available for use. When the mobile device has more than one physical adapter, the Switched Roaming feature would provide the user with the best interface to use.

The proposed Adapter Switching component consists of a Network Device Interface Specification (NDIS) 5.0-compliant intermediate driver, a mobility services client (also called the Adapter Switching service, which contains the Policy Manager), and a plug-in user interface (UI) component for the PROSet GUI Manager. The intermediate driver is responsible for Dynamic Host Configuration Protocol (DHCP) blocking (explained below), and has a mini DHCP client implementation. It also serves as a link monitor and informs the Policy Manager (PM) up in user space about link change events. The plug-in UI component for the PROSet UI allows the user to select a preferred adapter type, and it also displays the current status of the various network adapters present in the system.

The PM enforces a set of user-defined policies. It contains all the logic necessary to determine the most preferred state of the mobile device, given a dynamically changing network environment and a set of user-defined policies. The user will be able to set a preference level for the different network adapter types that are available. The PM will attempt to select an adapter that has the highest preference, as determined by the policy, and make that the primary adapter for all network connections on that mobile device. In the event that the preferred adapter type is unavailable, the PM will attempt to select another one from the list that it maintains.

The Adapter Switching component will be enabled only when all supported adapters are DHCP-enabled—if any adapter is assigned a static IP address, the Adapter Switching component will disable itself and appropriately notify the user. Normally, all adapters in a connected state on the mobile device will be able to obtain an IP address; the route-table on the Windows machine will contain routes using all the available adapters on that machine. If there is more than one adapter in the system that is available, and if the route-table has the same metric for all routes, then the particular adapter that will be chosen for data transfer will be indeterminate (an adapter for data transfer is chosen based on a combination of longest prefix match and route-table metrics). The Adapter Switching component's job is to deterministically present a single best interface for use, given that multiple interfaces are available.

This is achieved by allowing the TCP/IP stack to obtain an IP address for only one interface at any given time—and this will be the "active" interface. This means that the route-table will contain entries for only one interface, and there will no longer be any ambiguity when it comes to choosing an interface for data communication. The intermediate driver thus blocks DHCP traffic on all but the active interface, and prevents the TCP/IP stack from obtaining an address on an "inactive" interface. The interface to be made active is chosen by the PM, depending on network conditions and the user's policy, and this decision is communicated to the intermediate driver when state transitions occur.

The proposed Adapter Switching component will also support *ad hoc* 802.11 connections (in either static or dynamic addressing mode) concurrently with other adapters in infrastructure mode. This means that if a WLAN adapter is set to be in ad hoc mode, it can be used for communication at the same time another adapter (in non-ad hoc mode) is used, without leading to a loss in the deterministic behavior that Adapter Switching provides.

Power Conservation

Another very useful feature of the Adapter Switching software is its ability to help save battery life by working in conjunction with Wireless PROSet.

Consider a scenario where a user Joe has prioritized wired adapters over WLAN adapters. When Joe is working at his desk, he is connected using his wired adapter; since WLAN deployment in his company is ubiquitous, he is also connected to the WLAN. When the Adapter Switching software is enabled, it detects the availability of both types of connections; however, since Joe has indicated his preference for the wired adapter, the Adapter Switching software will only allow his wired adapter to obtain an IP address and be used for communication. Since the WLAN card is essentially not used at this point, the Adapter Switching software informs Wireless PROSet that this adapter can be turned off. When Wireless PROSet receives this indication, it can potentially turn off the radio on the WLAN card and therefore save precious battery life.

When Joe unplugs his laptop and walks to a conference room, the Adapter Switching software will detect this change; since the preferred (wired) connection is no longer available, the software will request Wireless PROSet to activate the WLAN card. Once that happens, Adapter Switching will ensure that this adapter can acquire an IP address; traffic will now start flowing over this adapter.

This intelligent interaction between Adapter Switching and Wireless PROSet can thus help save power.

Summary

The Adapter Switching software thus offers the following benefits:

1. It offers a single best interface to use, depending on the user's preferences and the current state of the network. All traffic will be deterministically directed over this one interface, even when other less-preferred interfaces are available.
2. It offers one way to help conserve valuable battery power by intelligently turning off the WLAN adapter when it is not needed.
3. It increases the security of a system when a VPN is enabled. If a VPN client is enabled when multiple adapters are available, the VPN tunnel will be established over one particular interface, and all packets flowing through that particular interface will be encrypted. However, there is a possibility for unencrypted packets (which might contain sensitive information) to flow over another interface that is not protected by the VPN tunnel. The Adapter Switching software prohibits this from happening by ensuring that there is only one interface that is active at any given time.
4. It supports "split-tunneling" within an interface when a VPN tunnel is active, if this feature is also supported by the VPN client. This will enable access to local resources (such as a local printer) while in a VPN session.
5. It supports communication over a WLAN adapter in ad hoc mode while simultaneously supporting communication over the preferred interface (which is in non-ad hoc mode).

AUTOMATIC VPN INVOCATION

The Problem

WLAN deployment is quickly gaining traction, especially in enterprise environments, where employees expect to be able to move between buildings and conference rooms and be constantly connected to the network. The first question that will be asked by an IT administrator before she starts to deploy a WLAN in her enterprise will be about its security features. It's a well-known fact that the best currently available mechanism, Wired Equivalent Privacy (WEP), is broken. Other more secure mechanisms (such as 802.11i) are not yet standardized, and interim solutions (such as Wi-Fi Protected Access (WPA)) are not yet widely deployed. As a result, most enterprises require their users to protect all WLAN traffic with a VPN connection. This ensures that a person with malicious intent sitting in the parking lot of the company campus cannot gain access to the enterprise Intranet via the WLAN. It also protects (by encrypting) all user traffic that flows over the air, thus ensuring that the same person in the parking lot cannot just snoop the air for tasty tidbits of information.

While this is great from a security perspective, it is a virtual nightmare from a user's standpoint. When a user wants to use his wireless connection, for example at a conference room at work, he first has to figure out if his wireless adapter is connected to the network, and then remember to launch his VPN client and connect back into the enterprise before he can be "connected." When he then walks back to his desk and plugs in his wired card (because he likes a higher speed connection when one is available), he needs to remember that his VPN is currently running over the wireless card and that the wired card will be inaccessible to applications until the VPN client is turned off. And then when he needs to run to a meeting a few minutes later, he has to remember to launch the VPN client yet again to connect using the wireless network, and the whole process starts again.

Our Proposed Solution

In order to spare the user from this tedious and error-prone process, we propose an auto-launch feature that enables a specific VPN client to be launched at the right time; and not only that, the VPN client will also be torn down at an appropriate time.

The user will be allowed to specify all the information necessary to configure and enable the auto-launch feature; this information will be tied to a WLAN profile, as described in the following section.

The auto-launch feature can be supported both when Adapter Switching is enabled and when it is disabled. When a wireless profile with VPN configured is applied and the wireless card is currently "active," Wireless PROSet will automatically launch the specified VPN client (with a particular VPN tunnel, if configured). Subsequently, when a higher-preferred adapter becomes available, or the WLAN adapter loses connectivity, the VPN tunnel will be proactively torn down.

This feature thus enables the user to enjoy the freedom that a wireless connection provides, without having to be bogged down with all the details required to set up and configure that connection.

VPN Invocation Using Profiles

With this proposed solution, Wireless PROSet can support the concept of automated VPN connectivity via profiles. While setting up a WLAN profile, the user will have the option of selecting a VPN client (and a VPN profile, if supported by the VPN client) to launch when the profile is applied. For example, the user can create a profile called "Office," and select a VPN client to be launched when the profile is applied. Each time this profile is subsequently applied, the software will automatically launch the VPN client and connect to the VPN gateway using a VPN profile, if one was specified. The software therefore behaves in a proactive manner, and attempts to do as much for the user as possible, in terms of enabling the user to quickly and seamlessly connect to a preferred network.

The Adapter Switching component can automatically detect the presence of supported VPN clients, using a combination of two mechanisms: searching for the VPN clients and having the clients register themselves with our component using "VPN adapters" (described below). Once the Adapter Switching component detects a client(s), it will present that to the user via the Profile Wizard when the user creates a profile. The Profile Wizard then stores this information as part of its WLAN profiles; as soon as a profile is applied, the Wireless Configuration service will query its database to determine if the profile applied contained VPN-related information. If it detects that a VPN client needs to be launched, it will pass all necessary information to the Adapter Switching component that will then attempt to actually launch the client.

An Extensible Solution

Due to the wide variety of VPN client and gateway implementations, the architecture cannot be designed to work with all VPN solutions; this proposal for Quick Connect and Switched Roaming has been investigated with a select set of third-party VPN implementations. However, we propose an extensible mechanism that will enable other VPN implementations, if they meet certain requirements, to work with this architecture.

VPN clients from different vendors have their own implementations and expose different interfaces. VPN clients also have their own unique profile formats that are used to store all tunnel setup and configuration information. Furthermore, legacy versions of certain VPN clients allow only the client to be launched, while more recent versions allow the client to be launched along with a specific tunnel. In addition, these clients may also allow querying of their internal state to determine the status of a tunnel in place, and so on. In order to provide a common framework across these varied implementations, we have designed what is termed a "VPN adapter." This adapter implements a uniform interface that we have defined, and it essentially acts as an abstraction layer between our software and the actual VPN client. Using this adapter, we will be able to instruct the VPN client as needed, without having to worry about proprietary interfaces for particular VPN clients. We have designed built-in adapters for a select set of VPN clients; if a new client wants to co-exist with our software, and assuming it has passed the other requirements for co-existence (such as co-existing with our intermediate driver), all that it needs to do is implement this VPN adapter and register itself on the mobile device. Our software will then dynamically detect its presence and be able to utilize that VPN client.

CO-EXISTENCE WITH "SMART CLIENTS"

With the proliferation of WLANs come hotspots. These hotspots are public places where access to the Internet is provided through WLANs. Several wireless ISP's (WISPs) provide internet access at these hotspots. In order to gain access to these hotspot networks, the user first needs to be authenticated; the particular authentication mechanisms employed are sometimes proprietary methods. As a consequence, hotspot operators usually require users to install and use their own specialized client software for Authentication, Authorization, and Accounting (AAA). These specialized clients also help the user to look for and connect to the wireless networks at the hotspots.

A user with subscriptions to several WISPs could therefore potentially have a number of "smart clients" installed and active on her laptop. All of these clients, if active simultaneously, will be trying to search for and connect to networks that they each know about. This can lead to software co-existence issues and in-deterministic behavior, eventually resulting in a good amount of confusion on the user's part. One possible solution is to have Intel's Wireless PROSet software communicate with these various smart clients using a pre-defined mechanism. This solution is obviously not very extensible, and it would not be interoperable with clients that are already proliferating the market. Furthermore, this solution will work only with clients from vendors or operators that have chosen to embrace and implement this pre-defined mechanism. To circumvent these problems, Wireless PROSet employs a generic mechanism that is independent of smart clients from other vendors.

In the Windows* Operating System, applications communicate with the WLAN driver using either a standard OS-defined interface or some proprietary interface. While the Wireless PROSet software can use a combination of both standard and proprietary mechanisms, all hotspot and other AAA smart clients can only communicate with the driver using the OS-specified standard interface.

As mentioned above, Wireless PROSet communicates with the driver using a proprietary interface. When the OS boots up and Wireless PROSet "registers" with the driver, the driver monitors for a subset of "set-able" OS-defined wireless commands as part of the standard interface exposed by the driver. When a AAA smart client is active, the driver receives these standard commands from the smart client and deduces that these commands are from a smart client, and not from Wireless PROSet. It then notifies Wireless PROSet which in turn notifies the user that another smart client is also active. Depending on the action taken by the user, certain network connection-related features will get disabled in Wireless PROSet if the AAA client continues to remain active. At the end of this operation, only one client—either the Intel Wireless PROSet software or the smart client—will be guaranteed to be active and in control. Thus, this solution does not involve any direct client-client communication, but relies on information deduced from the client platform.

OVERALL SYSTEM ARCHITECTURE

Figure 1 below provides a conceptual overview of the system architecture as described in this paper.

*Other brands and names are the property of their respective owners.

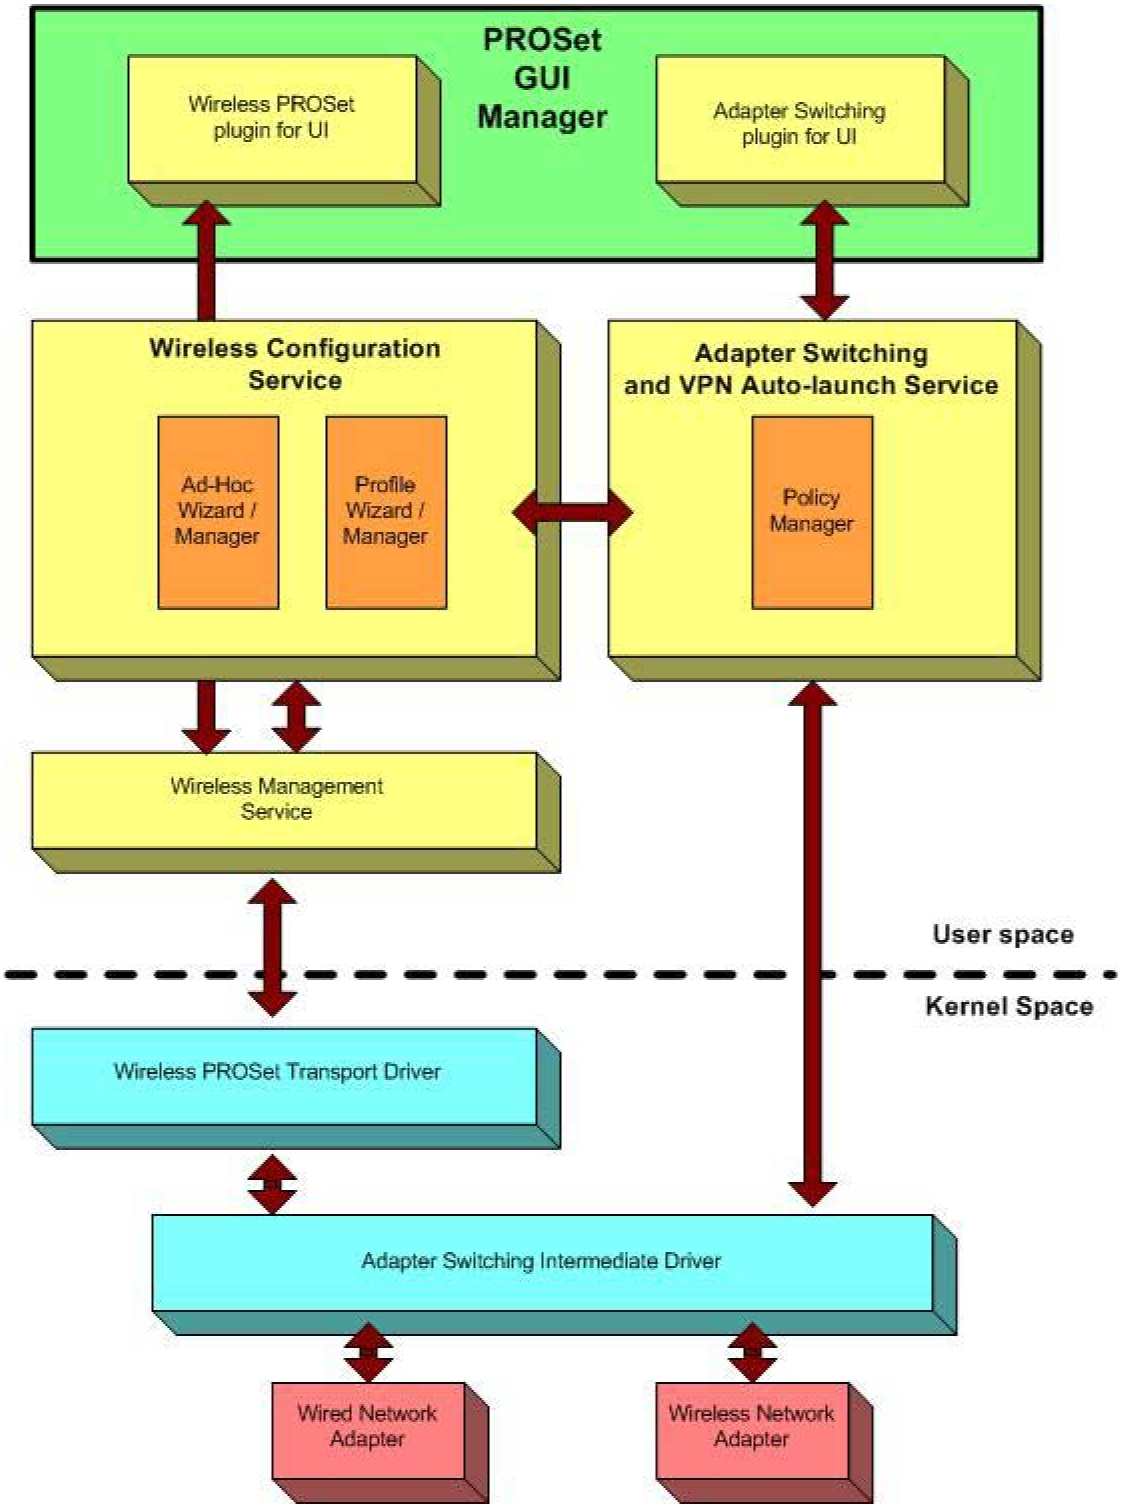
PROSet GUI Manager
Wireless PROSet plugin for UI
Adapter Switching plugin for UI
Wireless Configuration Service
Ad-Hoc Wizard / Manager
Profile Wizard / Manager
Adapter Switching and VPN Auto-launch Service
Policy Manager
Wireless Management Service
User space
Kernel Space
Wireless PROSet Transport Driver
Adapter Switching Intermediate Driver
Wired Network Adapter
Wireless Network Adapter

Figure 1: Overall system architecture

WALK-THROUGH OF A COMMON USAGE SCENARIO

Let's consider a common scenario, and see how all these components fit together. To make this a little more descriptive, let us picture a mobile user, Jane, who has a laptop with two network adapters. One network adapter is wired, and the other one is a built-in dual band wireless adapter, that can function as either an 802.11a or 802.11b adapter. The Intel PROSet software with the proposed features described in this paper is also running on Jane's laptop, to manage her network adapters.

Jane has used PROSet to configure two wireless profiles for her laptop—one profile (ProfileA) is for the 802.11a access points (APs) at her office, and the other (ProfileB) for the 802.11b APs (profile descriptions are illustrative only). Both these profiles have a Virtual Private Network (VPN) auto-launch setting associated with them, as the APs are deployed outside the company's Intranet. Using the Adapter Switching plug-in, Jane has given the wired adapter a higher preference than the wireless adapter.

When Jane walks in to work early Monday morning and powers up her laptop, the PROSet mechanisms are put into motion.

1. The Adapter Switching service detects both the wired and wireless adapters. It first verifies that all adapters are supported by the software. Once it detects that the wired adapter is plugged into an available link, it attempts to obtain a Dynamic Host Configuration Protocol (DHCP) address for the adapter. Simultaneously, the Wireless PROSet component also tries to associate with an available AP.
2. Once the Adapter Switching service verifies that the wired adapter can obtain a DCHP address, and since the wired adapter is the preferred adapter, it makes that the "active" adapter. What this means is that it will ask the operating system (OS) to renew the IP address on the active adapter, and release the IP address on all other adapters. The Adapter Switching component also notifies the Wireless PROSet component of this decision.
3. In the meantime, Wireless PROSet associated the wireless adapter with an AP matching ProfileA. The Adapter Switching service was able to internally verify that the adapter was able to acquire a DHCP address; however, since the wired adapter had a higher preference, it was the one that was made active.

At this point, Jane will be using her wired adapter for all network connections. After a while, she leaves her desk to attend a meeting in a conference room. She unplugs the wired adapter and takes her laptop to the conference room.

1. The Adapter Switching service detects that the wired adapter was unplugged, and it updates its state to register that there is currently no active adapter.
2. It then remembers that the wireless adapter is available for use. It makes the wireless adapter active and notifies the Wireless PROSet component of this.
3. The Wireless PROSet component then looks up its profile table and detects that there is a VPN client associated with this profile. It instructs the Adapter Switching service to establish a VPN tunnel and passes it the necessary information.
4. As soon as the Adapter Switching service successfully obtains an IP address for the wireless adapter, it attempts to create a VPN tunnel based on information passed by Wireless Proset in Step 3. It uses the VPN auto-launch component to invoke the VPN client with pre-configured parameters. Jane is required to enter her credentials as part of the VPN connection setup process.

Jane will now be using her wireless adapter, associated using ProfileA, for her network connections. Once the meeting is over, Jane plans on going down to the cafeteria to continue discussions with a colleague. As she takes the elevator down and walks into the cafeteria, she walks out of range of the first AP and into range of a different one.

1. The Profile Switching service within the Wireless PROSet component detects the change and uses the scan-list to determine the APs that are currently in range. It then applies the best profile it can, which happens to be ProfileB.
2. The Adapter Switching service detects the changes in the link state of the wireless adapter. As soon as the adapter disconnects from the first AP, it notifies the Wireless PROSet component, which in turn asks the VPN Manager component within the Adapter Switching service to shut down the VPN tunnel.
3. Once the wireless adapter associates with the new AP, the Adapter Switching service detects a new link, and once again it attempts to make this adapter active. Once that process is completed, it notifies Wireless PROSet. This, in turn, causes another VPN tunnel to be initiated, similar to the process described above.

Jane can now use her wireless link once again without having to determine whether she needs a VPN tunnel, or whether she has to manually initiate the VPN connection process. Most of the link state changes and adapter switches were handled transparently, and she can continue her work without having to manually connect to networks or subsequently remember to turn VPN clients on and off.

THE CHALLENGES AHEAD

In the last two years the markets for wireless mobile data have gained momentum and are moving from early adopters to mainstream users. Wireless Local Area Networks (WLANs), General Packet Radio Service (GPRS), and Bluetooth* technology, respectively, are the most common mobile data technologies, with each technology addressing complementary mobile data connectivity needs. Generally speaking, WLANs provide access connectivity within buildings or hotspots; GPRS data cards or GPRS phones enabled with Bluetooth technology offer wireless connectivity everywhere else.

As the adoption of these technologies increases so does the need for these technologies to seamlessly work together. The long-term goal is for the end-user to be oblivious of the data connection type, and still have access to the best network, without having to switch connections, restart applications, or reboot mobile computers. The software solution described in this paper is a first step in addressing this need by providing a mobile client that communicates with industry-standard infrastructure solutions.

Accomplishing the above would provide a mobile user with a truly seamless roaming experience using industry standards.

ACKNOWLEDGMENTS

The authors express their gratitude and appreciation to the entire management, development, and software quality teams that worked on the Wireless PROSet and Adapter Switching product. We thank Prakash Iyer and Marc Meylemans for reviewing this paper and enhancing its content. Finally, we thank Marian Lacey for doing a great job editing this paper.

AUTHORS' BIOGRAPHIES

Allan Chin is a staff software engineer in the Wireless Product Development division at Intel Corporation, and is currently based in San Diego, California. Allan has over 20 years of software engineering experience in a wide variety of fields spanning real-time, embedded, navigation systems to user-based Windows applications. Allan received his Masters degree in Computer Science from Stevens Institute of Technology in 1984. He holds a Bachelors degree in Electrical Engineering from the University of Delaware. His e-mail is allan.chin@intel.com.

Ajay Gupta is a staff software engineer in Intel's Wireless Product Development (WPD) division at San Diego, California. His primary interests include network protocols, networked multimedia, and image processing. Ajay joined Intel in 1996. Prior to joining WPD, he worked on video conferencing and network load balancing. He has a Masters degree in Computer Science. His e-mail is ajay.g.gupta@intel.com.

Ranjit Narjala is a network software engineer in Intel's Corporate Technology Group at Portland, Oregon. His professional interests include wireless technologies, networking, and mobility protocols. Ranjit received his Masters degree in Information Networking from Carnegie Mellon University in 2000. He also holds a Bachelors degree in Computer Science and Engineering.. His e-mail is ranjit.s.narjala@intel.com.

Venkata Vallabhu is a Senior Software Engineer in Intel's Wireless Product Development division at San Diego, California. He is responsible for the design and development of the components in wireless PROSet software at Intel. He has over six years of industry experience in software development in various Wireless Technology and Systems Engineering positions. Venkat holds a B.Tech degree in Electronics and Communications Engineering from Jawaharlal Nehru Technological University (India). His e-mail is. venkata.r.vallabhu@intel.com.

* Bluetooth is a trademark owned by its proprietor and used by Intel under license.

For further information visit:

developer.intel.com/technology/itj/index.htm