



Wi-Fi™ (802.11b) and Bluetooth™:

An Examination of Coexistence Approaches

Abstract

This paper analyzes different approaches to resolving the interference problems between the Wi-Fi™ and Bluetooth™ wireless technologies. This analysis explores the strengths and weaknesses of these interference mitigation approaches, and goes on to explain what is necessary for achieving satisfactory combination performance and true “Coexistence without Compromise”™. The contents are based on Mobilian Corporation’s coexistence research and development work, including a thorough analysis of the problems and various experiments to understand the interference issue.

In investigating different approaches to interference mitigation, this paper gives technical data and uses common wireless technology terms. The information presented is targeted to readers who have a basic understanding of wireless networking. We recommend that readers without this understanding read Mobilian’s first white paper: *Wi-Fi (802.11b) and Bluetooth Simultaneous Operation: Characterizing the Problem* (www.mobilian.com).

April 11, 2001

Disclaimer and Copyright

Windows Hardware Engineering Conference

Author’s Disclaimer and Copyright: *Mobilian™, TrueRadio™, Sim-OP™, TrueConnectivity™, and Coexistence without Compromise™* are all trademarks of Mobilian Corporation.

WinHEC Sponsors’ Disclaimer: *The contents of this document have not been authored or confirmed by Microsoft or the WinHEC conference co-sponsors (hereinafter “WinHEC Sponsors”). Accordingly, the information contained in this document does not necessarily represent the views of the WinHEC Sponsors and the WinHEC Sponsors cannot make any representation concerning its accuracy. THE WinHEC SPONSORS MAKE NO WARRANTIES, EXPRESS OR IMPLIED, WITH RESPECT TO THIS INFORMATION.*

Microsoft, DirectX, MS-DOS, Win32, Win64, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Other product and company names mentioned herein may be the trademarks of their respective owners.

Table of Contents

1.0	Executive Summary	3
2.0	Background	4
2.1	Technical Background	5
2.1.1	Bluetooth™ Wireless Personal Area Networking (WPAN)	5
2.1.2	802.11b Wireless Local Area Networking (Wi-Fi™)	6
2.1.3	Wi-Fi™ / Bluetooth™ Interaction and Interference	6
3.0	Interference Mitigation Approaches	8
3.1	Collocation without Coexistence Mechanism	9
3.1.1	Overview	9
3.1.2	Analysis	9
3.2	Driver-level (Modal) Switching Between Wi-Fi and Bluetooth	10
3.2.1	Overview	10
3.2.2	Analysis – Dual-mode Radio Switching	11
3.2.2.1	Dual-mode Radio Switching	11
3.2.2.2	Leaving the Network without Signaling	12
3.2.2.3	Leaving the Network with Signaling	12
3.2.3	Analysis – Driver-level Switching	12
3.2.3.1	Throughput and Time-delay Concerns	12
3.2.3.2	Impacts of Bluetooth Polling Activities	14
3.3	Adaptive Hopping	14
3.3.1	Overview	14
3.3.2	Analysis	15
3.3.2.1	Adaptive Hopping as Optional Profile (Operational Mode)	15
3.3.2.2	Adaptive Hopper Must Accurately Sense and Respond to Interferers	15
3.3.2.2.1	Bluetooth™ Difficulty in Detecting Wi-Fi™ Signal	16
3.3.2.2.2	Congested Wireless Environments are Particularly Troublesome	17
3.3.2.3	Adjacent-Channel Noise	19
3.3.2.4	Number of Channels	19
3.4	MAC-level Switching	19
3.4.1	Overview	19
3.4.2	Analysis	19
3.5	Simultaneous Operation	20
3.5.1	Overview	20
3.5.2	Analysis	22
4.0	Summary	22
5.0	Appendix 1 – In-band versus Out-of-band Noise	24
5.1	Signals and Noise	24
5.1.1	Types of Noise	24
5.2	Bluetooth and Wi-Fi Interference Cases	26
6.0	Appendix 2 – Path Loss Models Employed	28
References	Error! Bookmark not defined.	

Table of Figures

Figure 1 – Performance Hierarchy of Coexistence Approaches for Collocated Wi-Fi & Bluetooth	9
Figure 2 – Geometry of Measurement and Simulation Environment	10
Figure 3 – Measurement of Wi-Fi Throughput in the Presence of Collocated Bluetooth	10
Figure 4 – Conceptual Wireless System Diagram	11
Figure 5 – Basic Geometry of Bluetooth and Wi-Fi Penetrated Corporation	17
Figure 6 – Likely Location of Two Adaptive Hoppers	18
Figure 7 – Simultaneous Operation Covers Entire Conceptual Wireless System Diagram	21
Figure 8 – Ganymede Chariot Graph of Mobilian Corporation's TrueRadio™ Demonstration	22
Figure 9 – Mobilian's TrueRadio™ Performance in Collocated Scenario	23
Figure 10 – In-Band versus Out-Of-Band Noise	25
Figure 11 – White Noise and Colored Noise are Very Different	25
Figure 12 – Typical Transmit Mask	26
Figure 13 – Path Loss as a Function of Distance, Indoor, 2.4-GHz ISM Band	28

Executive Summary

Wireless markets, from wide area networks, to local and personal area networks, are widely expected to be the significant market of the 21st Century. Investment capital is flowing to wireless companies worldwide, and market forecasts consistently project hundreds of millions of installed units. With this expansion comes increased opportunity for market innovation, and consequently, wireless penetration into the core fabric of our everyday lives.

This growth is spurred by increasing demand for maximum convenience and immediate access to desired information. It is facilitated by an unlicensed frequency spectrum, providing unlimited, free access to whomever wishes to build a wireless device capable of complying with regulatory standards. These forces are working together to create traffic and device density in the unlicensed frequencies, and consequently opportunities for interference between the protocols using those frequencies.

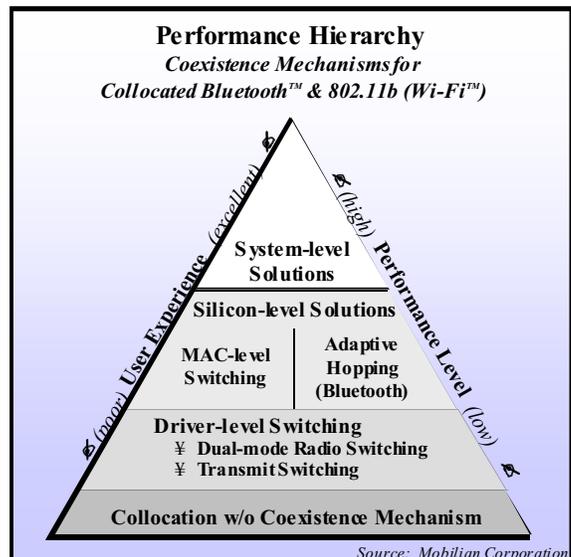
As these trends develop, the need for multiple wireless devices operating at the same time will increase, resulting in still greater potential for interference. That is why “simultaneous operation” is becoming an important topic of discussion in today’s market.

Simultaneous operation is the ability of different, fully standards-compliant wireless systems to operate simultaneously in any scenario, while experiencing minimal or no degradation in performance. This definition includes wireless devices that can give the user outstanding performance without a list of operational caveats. The device should “just work,” regardless of other devices within its operating environment.

Wireless local area networking (WLAN) and wireless personal area networking (WPAN) are two networks in particular for which simultaneous operation is growing in importance. WLAN / WPAN simultaneous operation will occur more and more frequently as users begin completing everyday tasks such as copying or printing a file from their WLAN PC while using a WPAN-enabled mouse,

keyboard, and speakers. Its frequency will continue to grow as personal communication devices and synchronization activities with PCs and networks grow, and it will gain even more importance as distributed applications take off – “the next big thing in software” – and WPAN devices must coexist with massive amounts of WLAN activity.

In all these scenarios, users will appreciate being able to use whatever wireless devices surround them, when they want to, and how they want to. Users will demand unhindered simultaneous operation and will resist adopting wireless devices as long as there are operational problems or perceived concerns.



With this certainty facing the market today, regulatory bodies, standards bodies, and industry participants are begun several approaches to achieving simultaneous operation, including:

1. Simple collocation (combo-card reference designs);
2. Approaches in the host software (driver-level switching and dual-mode radios)
3. Approaches in the MAC layer (MAC-level switching and adaptive hopping); and
4. System-level solutions covering the entire wireless sub-system, and incorporating the best aspects of many different approaches.

Each technique's strengths and weaknesses are explored in depth in the following pages. In summary, based on exploration and assessment of each technique's interference management ability, true, sustainable simultaneous operation can only be achieved by taking a system-level approach across the entire wireless sub-system. This allows the simultaneous operation solution to selectively use the best aspects of all the techniques, and therefore manage interference extremely well. This is the approach Mobilian Corporation has employed with its first product, TrueRadio™.

Background

The 2.4 GHz Industrial, Scientific, and Medical (ISM) band is poised for strong growth. Fueling this growth are two emerging wireless technologies: WPAN and WLAN. The WPAN category is led by a short-range wireless technology called Bluetooth™. Designed principally for cable replacement applications, most Bluetooth implementations support a range of roughly 10 meters, and throughput up to 721 Kbps for data or isochronous voice transmission. Bluetooth is ideal for applications such as wireless headsets, wireless synchronization of PDAs with PCs, and wireless PC peripherals such as printers, keyboards, or mice. Cahners In-Stat predicts shipments for Bluetooth devices will reach 800 million units annually by 2004 [CIS00a].

In the WLAN category, several technologies are competing for dominance; however, based on current market momentum, it appears that Wi-Fi™ (IEEE 802.11b) will prevail. Wi-Fi offers throughput up to 11 Mbps and covers a range of approximately 100 meters. With WLANs, applications such as shared Internet access, e-mail, and file sharing can be done in the home or office, resulting in new levels of freedom and flexibility. Cahners predicts WLAN shipments exceeding 38 million units annually in 2004, implying an installed base of nearly 95 million systems [CIS00b] by the same year.

“Coexistence,” the ability for multiple protocols to operate in the same frequency band without significant degradation to either's operation, has recently become a significant topic of analysis and discussion throughout the industry. This is due to several factors. Both protocols are expecting rapid growth, and because they both operate in the 2.4 GHz frequency band, the potential for interference between them is high. Also, WPAN and WLAN are complementary rather than competing technologies. Consequently, more and more usage models are being discovered in which it is desirable and necessary for both Bluetooth and Wi-Fi to operate simultaneously and in close proximity.

Technical Background

This section provides some high-level background on several key characteristics of the Bluetooth and Wi-Fi protocols. A deep understanding of these characteristics is necessary to fully investigate the merits of various approaches, but this high-level overview will provide a basic understanding. Further explanation of the two protocols' technical characteristics is provided in Mobilian's first white paper, *Wi-Fi (802.11b) and Bluetooth Simultaneous Operation: Characterizing the Problem* www.mobilian.com/whitepaper_frame.htm [MBLN01].

Bluetooth™ Wireless Personal Area Networking (WPAN)

Bluetooth is a WPAN protocol designed as a cable-replacement technology - low cost, modest speed, and short range (<10 meters). Bluetooth can support piconets of up to eight active devices, with a maximum of three synchronous-connection-oriented (SCO) links. SCO links are voice-oriented and designed to support real-time, isochronous applications such as cordless telephony or headsets. Bluetooth also supports asynchronous connection links (ACLs) used to exchange data in non-time-critical applications. The majority of Bluetooth devices transmit at a power level of 1 mW (0 dBm). The Bluetooth physical (or PHY) layer uses the frequency-hopping spread spectrum (FHSS) technique. Bluetooth hops at a rate of 1600 hops/sec and uses Gaussian frequency shift keying (GFSK) modulation.

When the Bluetooth technology establishes communication, it forms small networks, or piconets, of Bluetooth-enabled devices. Piconet topology consists of a single master and up to seven active slaves. In a single piconet environment, there can be only one Bluetooth device transmitting in any single time slot at any one time. Therefore, the master Bluetooth node of the piconet controls the piconet through a series of transmissions. When the master has information to transmit to the slaves, it does so. Otherwise, the master is constantly polling the slaves and listening for their responses¹. In short, for a slave to transmit data, it first must be "asked" to do so. The slave's responses can be either NULL for no information to transmit, or they can begin transmitting if they have information to transmit. This piconet management scheme avoids interference within the piconet and is standard for any device carrying the Bluetooth certification (i.e., complying with the Bluetooth specification). Understanding some aspects of the different approaches to interference mitigation, requires further investigation of the master/slave polling mentioned above. Due to the extremely rapid nature of the polling activity (hundreds of microseconds), the Bluetooth media access controller (MAC) controls the function at the MAC-level and thus, the data transferred in the process is not

¹ The Bluetooth specification does not dictate how often a master should poll a slave, nor does it provide for any preemptive transmission from the slave to inform the master that it has data to transmit; therefore, to maximize Bluetooth throughput, many typical current design practices call for the master to poll the slaves during every available transmit time slot (800 polls / second) while in an active piconet.

made available at the driver or host level. This will prove to be very significant, as is explained in later sections.

802.11b Wireless Local Area Networking (Wi-Fi™)

Like wired Ethernet, Wi-Fi supports true multipoint networking with such data types as broadcast, multicast, and unicast packets. Although standard practice is approximately one access point (AP) to every 10-20 stations (STA), the MAC address built into every device allows for a virtually unlimited number of devices to be active in a given network. These devices contend for access to the airwaves using a scheme called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). The Wi-Fi physical layer uses direct-sequence spread spectrum (DSSS) at four different data rates using various modulation techniques to communicate. The transmit power level can vary, but is typically between 30 and 100 mW (+15 to 20 dBm).

Wi-Fi™ / Bluetooth™ Interaction and Interference

Bluetooth and Wi-Fi share the same unlicensed 2.4 GHz ISM band that extends from 2.4 to 2.4835 GHz under US FCC regulations. This frequency band is free of tariffs under the ISM band rules defined in FCC Part 15.247 [FCC15.247]. However, systems in this band must operate under certain constraints that are supposed to enable multiple systems to coexist in time and place. FCC Part 15.247 specifies that a system can use one of two methods to transmit in this band: FHSS or DHSS.

FHSS is a technique in which a device transmits an energy burst in a narrow frequency band for a limited time before it hops to another. This hopping process is repeated rapidly across the entire frequency band in a pseudo-random fashion.

DSSS is a technique in which a device communicates by distributing its energy across a defined set of contiguous frequency bands without hopping.

Bluetooth is an FHSS technology with frequency channels 1 MHz in width and a hop rate of 1600 hops per second. Bluetooth dwells 625 μ sec in every frequency channel. In the United States and most of the world, Bluetooth uses 79 different 1 MHz frequency channels of the available 83.5 MHz in the 2.4 GHz ISM band.

Wi-Fi uses DSSS with a 22 MHz passband, and communicates with throughput up to 11 Mbps. A Wi-Fi system can use any of eleven² 22-MHz wide sub-channels across the available 83.5 MHz of the 2.4 GHz frequency band.

Because Bluetooth hops on 79 of the available 83.5 1-MHz channels, and Wi-Fi occupies 22 1-MHz channels within its passband, sharing between the two technologies is inevitable. Two wireless systems using the same frequency band will have a high propensity to interfere with each other.

² The 11 sub-channels available under US regulation allow for multiple variations of locations for 3 simultaneously operating Wi-Fi networks and associated passbands. A Wi-Fi passband typically spans a 22-MHz channel; therefore the 83.5 MHz available within the 2.4 GHz band can support three simultaneously operating, overlapping Wi-Fi networks (83.5 MHz - (3*22 MHz) = 17.5 MHz). Geographies outside of the US may support more or fewer than 11 selectable sub-channels.

In October of 2000, Mobilian Corporation published a white paper that explored this interference in great detail. The white paper, *Wi-Fi (802.11b) and Bluetooth Simultaneous Operation: Characterizing the Problem*, received wide acceptance by the industry as the definitive treatment of this issue. This current paper, on the other hand, builds on the previous work and therefore assumes a certain level of understanding of the coexistence issues.

However, for the basis of this paper, it is important to establish that Wi-Fi performance generally suffers more from Bluetooth activity than vice versa. The reasons for this are explained in great detail in the aforementioned white paper, but in summary, there are two main reasons:

- 1) First, the Wi-Fi MAC is an adaptation of the wired Ethernet MAC, and therefore uses carrier-sense before transmission (also known as “listen before talk”). Unlike wired Ethernet, the Wi-Fi MAC cannot detect collision, so Wi-Fi dictates that every received packet is acknowledged by an “acknowledgement” (ACK). If a station or access point transmits a packet and does not receive an ACK from its target recipient, it assumes a collision with another Wi-Fi transmission has occurred. To avoid additional Wi-Fi collisions, the station uses an exponential back-off algorithm (i.e., pauses a few micro-seconds) and transmits again. By using this mechanism among others, wired and wireless Ethernet work very efficiently in a homogenous environment. However, in an unpredictable and highly interference prone Bluetooth/Wi-Fi environment, this mechanism, and its associated back-off algorithms, result in repeated error correction without corresponding interference improvement, ultimately resulting in reduced Wi-Fi throughput.
- 2) Second, the Wi-Fi protocol does not typically move from its 22 MHz passband³. This renders it highly susceptible to collision with Bluetooth. Roughly, the probability that a standard Wi-Fi 1500 byte transmission will collide with a simultaneous Bluetooth transmission is 55%. This results from the fact that Wi-Fi requires approximately 1 to 1.5 milli-seconds to receive a 1500 byte packet at 11 Mbps. This allows Bluetooth to hop approximately 2 times (625 μ sec per hop / 1.5 milli-seconds). Each hop has a 1 in 79 chance of hitting a given channel, therefore 2 hops have a 2 in 79, or $\sim 1/40$, chance of hitting a given channel. With 22 channels occupied by the Wi-Fi network, this raises the probability to $\sim 22/40$ or $\sim 55\%$.

This performance degradation occurs at any one of three levels in descending order of severity.

- 1) The most pronounced negative effect occurs when a Bluetooth device is collocated with a Wi-Fi device, as is the case in a combination card or notebook PC with both Wi-Fi and Bluetooth functionality.
- 2) The effects are slightly less severe when the transmitting Bluetooth device is located within the same piconet as a collocated Bluetooth

³ Wi-Fi does have “channel agility” functionality; however, it is seldom used and even if it is employed, due to its relatively slow movement between channels, it is practically ineffective in avoiding the extremely rapid BT hopping pattern.

and typically within 1 to 1_ meters from the collocated Bluetooth/Wi-Fi device.

- 3) The least severe effects occur when the interfering Bluetooth is outside the collocated Bluetooth’s piconet and more than 2 meters from the collocated device.

Additional factors can either improve or worsen the negative effects outlined above. One the most important is in-band and out-of-band communication of the two protocols⁴. Table 1 below gives an overview of the different scenarios and their relative severity.

		Bluetooth Tx		Bluetooth Rx	
		In-band	Out-of-band	In-band	Out-of-band
Wi-Fi™	Tx	No Conflict	No Conflict	Strong Interference	Moderate Interference
802.11b	Rx	Strong Interference	Moderate Interference	Strong ⁵ Interference	Moderate Interference

Source: Mobilian Corporation

Table 1: The Interference Cases for Bluetooth and Wi-Fi⁵

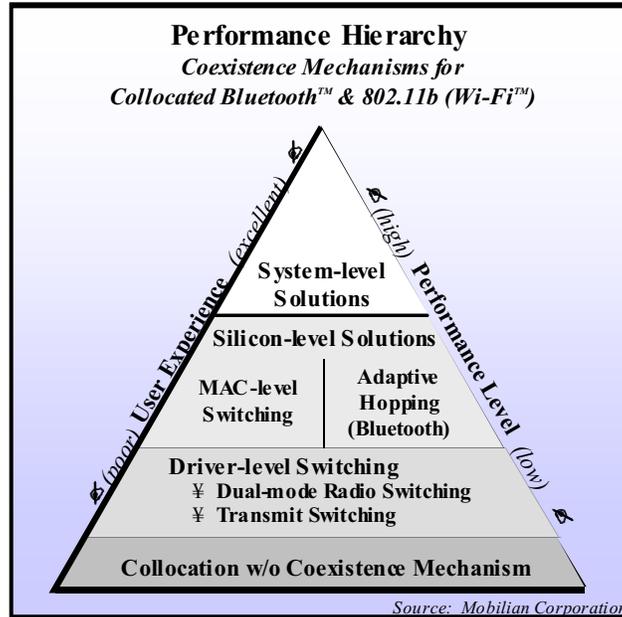
Interference Mitigation Approaches

As a result of the potentially negative impacts of collocated Wi-Fi and Bluetooth devices, many companies have begun researching and developing solutions for coexistence. Potential approaches include:

- Simple device collocation with no coexistence mechanisms;
- Restricted or adaptive band hopping for Bluetooth devices;
- Switching between the two protocols; and
- System-level approaches covering the entire wireless sub-system and many of the above techniques.

⁴ In-band refers to simultaneous operation in the same frequency channel. Out-of-band refers to simultaneous operation in two separate channels. This is further explained in the first white paper and in the appendix of this white paper, “6.0 Appendix – In-band versus Out-of-band Noise”. This appendix is an excerpt from Mobilian’s first white paper.

⁵ Collocated receivers is not an interference issue. However, simultaneous reception implies some degree of simultaneous transmission by external wireless systems. In the case of collocated 802.11b and Bluetooth systems, transmissions (which the collocated Bluetooth is trying to receive) from nearby Bluetooth nodes (located within 2 meters), can significantly affect 802.11b’s ability to receive.



**Figure 1 – Performance Hierarchy of Coexistence Approaches
for Collocated Wi-Fi & Bluetooth**

Each of these approaches is explored in the following pages and can be categorized into the performance and user experience hierarchy shown in Figure 1. The performance hierarchy could change dependent on the operating characteristics of the particular environment. In some scenarios, MAC-level switching may manage interference more effectively than adaptive hopping, and vice versa. The same can be said of driver-level switching and its various implementations. However, system-level solutions, providing simultaneous operation through a combination of the most appropriate aspects of each technique, will most consistently appear at the pinnacle of both performance, and user experience.

Collocation without Coexistence Mechanism

Overview

This approach simply entails collocating the two wireless devices in a single form factor without any attempt to avoid the potential interference (e.g., PC NIC reference design).

Analysis

Collocating Bluetooth and Wi-Fi without using any coexistence mitigation techniques increases the likelihood of significant interference. The coexistence issues associated with it are fundamental to the interference problem, which we have explored extensively in our first white paper. Performance is likely to be significantly degraded for both protocols in this scenario. Figure 3 shows both measured and simulated effects of this approach in the single-user network

configuration shown in Figure 2. The first white paper provides extensive details of both the scenario below and simulation details.

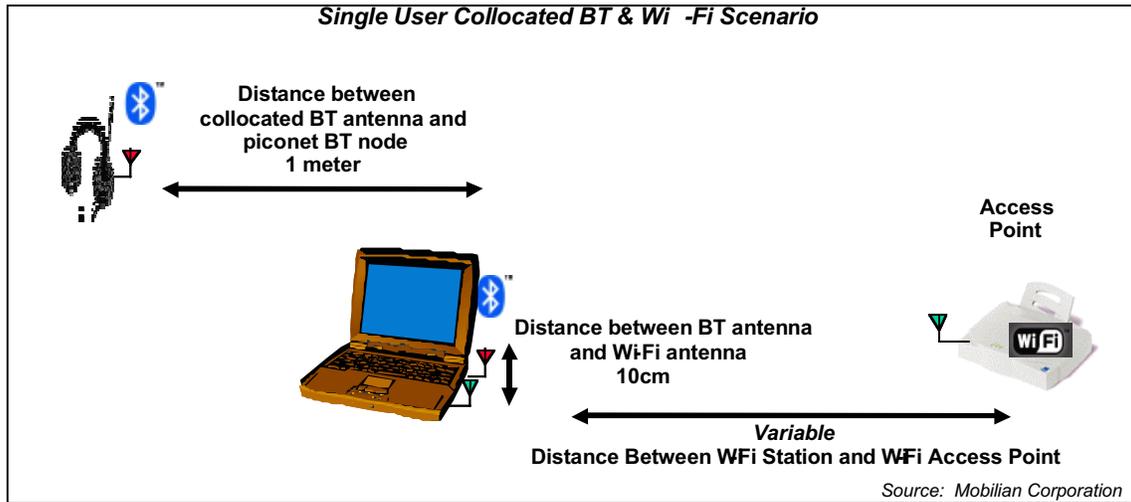


Figure 2 – Geometry of Measurement and Simulation Environment

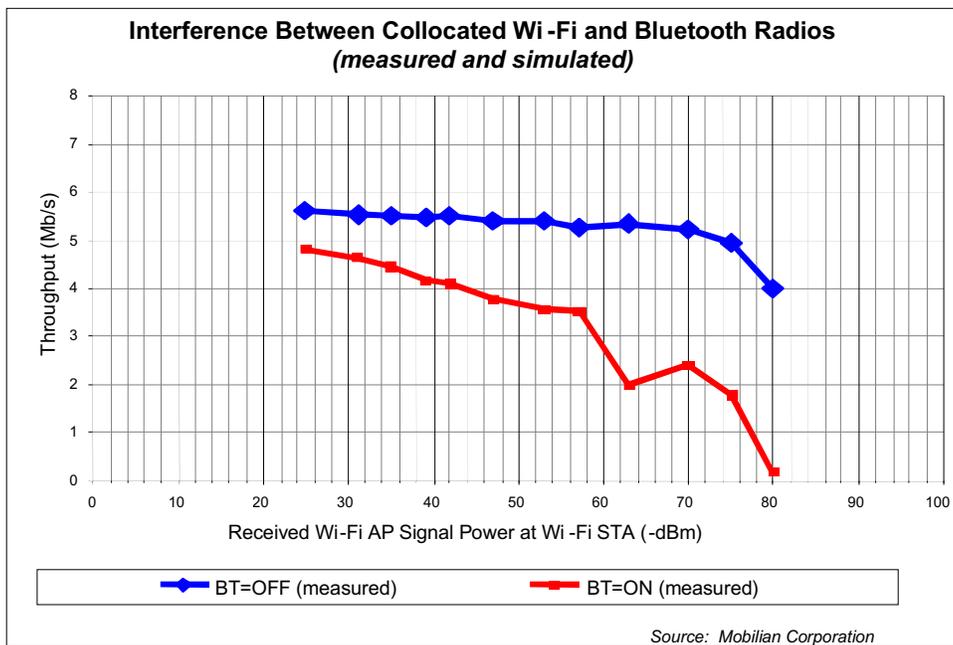


Figure 3 – Measurement of Wi-Fi Throughput in the Presence of Collocated Bluetooth

Driver-level (Modal) Switching Between Wi-Fi and Bluetooth

Overview

Driver-level switching is a time-division approach, essentially dividing the operational periods for each radio, and has many possible implementations.

Each different driver-level implementation generally adheres to the characteristics described in the analysis section below; however, dual-mode radio switching has several slight differences that are explored independently. The various forms of driver-level switching solutions include:

- 1) *Dual-mode radio switching* – The system shuts off one of the two radios completely when the other is operational (e.g., placing Bluetooth in park/hold mode or Wi-Fi in power-save mode). This is accomplished either through signaling or no-signaling approaches.
- 2) *Driver-level switching* – This includes several types of techniques that are all controlled at the driver level: User-dependent switching, discriminatory switching, successful-transmission switching, statistical switching, and time-delay switching.

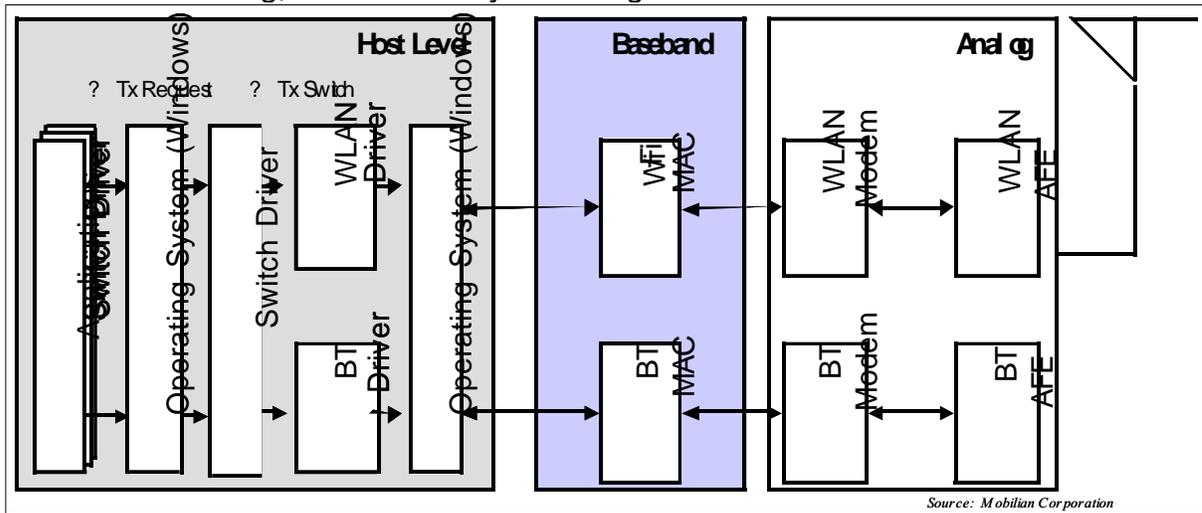


Figure 4 – Conceptual Wireless System Diagram

To effectively address driver-level switching, it is important to understand how drivers work with wireless radios. Figure 4 provides a high-level, conceptual graphic of a wireless system. Referring to the numbers in Figure 4, (◆) host applications initiate a request to transmit data. This request first travels to the operating system (e.g., Windows), then to the driver, (❖) which passes the message through the operating system again thus allowing the data to be transmitted via the wireless system. In a driver-level switching approach, the switch driver monitors these application requests to ensure no transmission collides with another.

Analysis – Dual-mode Radio Switching

Dual-mode Radio Switching

A dual-mode radio switching approach involves shutting one of the two radios off completely when the other is operating (e.g., placing Bluetooth in park/hold mode or Wi-Fi in power-save mode). The radios are never operating simultaneously, and therefore never attempting to simultaneously receive or transmit. This

approach avoids interference from Bluetooth polling, an important technical difficulty with other driver-level switching approaches explained in section 3.2.3.2. Dual-mode radio switching can be accomplished either by simply stopping operation of one of the radios with no indication to other devices in the network, or by first signaling that one device is about to be suspended and then stopping operations.

Leaving the Network without Signaling

In a normal operating scenario, radios do not go into these “sleep” modes commonly used for saving power, unless they are not actively participating in the network. When a radio is suspended without signaling to other partner devices, it cannot respond to transmissions from other network nodes. This lack of information, usually in the form of ACKs, results in reduced throughput and wasted bandwidth as the AP repeatedly sends data and executes interference mitigation techniques (discussed in section 2.1.2).

Leaving the Network with Signaling

By notifying other nodes in the network that a device is being suspended, problems with unacknowledged transmissions can be avoided. Thus while throughput remains a concern because modal approaches inherently reduce system up-time, this approach will lessen interference impacts. It will not, however, convince users that they are experiencing simultaneous operation, for the following reasons.

The time involved in stopping and reinitializing radio operation is time lost to the overall performance of the system. In most Bluetooth radios, this time is approximately 9 milli-seconds for 7 nodes in the piconet. For Wi-Fi devices, the amount of time varies according the network configuration, but is generally very short (1-3 milli-seconds). While this is certainly not onerous, the cumulative impact of repeated cycles and associated time delays could easily be noticeable to the end-user.

Neither of these solutions, signaling or no signaling, will manage Bluetooth synchronous-connection-oriented links (SCO), or voice links. Bluetooth SCO links are very timing-sensitive and cannot be interrupted by Wi-Fi activity under this approach. This could potentially lead to poor user experiences when the user is attempting to talk on a Bluetooth headset and simultaneously search for files on the server or intranet using Wi-Fi.

Analysis – Driver-level Switching

Throughput and Time-delay Concerns

An intuitive problem with all modal (on/off) switching is the impact to the protocols' throughput. If a radio is suspended, it is not transmitting or receiving, and therefore the potential throughput is degraded. While this can be significant, it varies according to the differences in implementation. Therefore, we do not

address the issue here. Rather, we investigate the more important aspects of the overall approach.

Because driver-level transmit switching occurs at the driver level, and because the transaction time for a driver-level switch to occur is unreliable and lengthy, avoiding collisions with incoming packets is very difficult. The resulting transmission of one protocol during reception of the other causes loss of received packets, interference, and potential user difficulties. This is caused by the driver's dependence on the host operating system, which is generally non-deterministic in its response time (i.e., non-real-time).

Reception of a standard 1500 byte Wi-Fi packet takes approximately 1 to 1.5 milli-seconds. The time required for information to transmit from the baseband to the driver – such as “there is a packet being received – do not transmit” – and for the corresponding driver activity to complete, can be anywhere from 100 micro-seconds to 2 or 3 seconds, or even longer. This is caused by the variable latency inherent in non-deterministic (non-real-time) host operating systems such as Windows, Linux, and Unix.

Host operating systems have variable latency because of the many background activities occurring during normal operation. As the interrupt from the baseband is received by the operating system, it is queued behind other interrupts and requests from other functions. This queue could range from very short to very long in terms of time required for the operating system to process the baseband interrupt. Because of this varied latency, when the operating system processes the request and passes it to the driver, the driver is not able to gauge a proper response. It doesn't “know” if the baseband request was sent 1 micro-second ago, 1 milli-second ago, or 1 second ago. This represents a huge gulf of missing information to Wi-Fi and Bluetooth, which both operate in micro-second intervals. For this reason, basebands do not currently perform this activity and a driver-level approach will potentially transmit at the same time the wireless system is receiving. As we illustrated before in Table 1 and again in Table 2 below, this scenario, one radio transmitting while the other is receiving, causes significant interference.

		Bluetooth Tx		Bluetooth Rx	
		In-band	Out-of-band	In-band	Out-of-band
Wi-Fi™	Tx	No Conflict	No Conflict	Strong Interference	Moderate Interference
802.11b	Rx	Strong Interference	Moderate Interference	Strong ⁶ Interference	Moderate Interference

Source: Mobillian Corporation

Table 2: The Interference Cases for Bluetooth and Wi-Fi⁶

Impacts of Bluetooth Polling Activities

As we indicated in our discussion of Bluetooth polling functionality in Section 2.1.1, in most implementations of Bluetooth, in an active piconet, the master Bluetooth node will continuously poll the slaves. “Continuously” means in every available transmit slot or 800 times per second (1 slot request for information, 1 slot opportunity to respond).

As also explained, polling activities are controlled at the Bluetooth MAC layer and don’t reach the driver-level; therefore, polling activities cannot be controlled / switched by the driver. Again, this creates significant interference, because Bluetooth will be continuously transmitting while Wi-Fi is attempting to receive. The effect of Bluetooth polling activities on Wi-Fi performance generates significant interference roughly equivalent to that in collocated Wi-Fi and BT radios with no coexistence mechanism as in Figure 3.

Adaptive Hopping

Overview

Recently (11/13/00), a group of companies petitioned the FCC requesting the initial report and order (R&O) for Wideband Frequency Hopping (WBFH), also known as ET Docket 99-231, be amended or reconsidered to allow Bluetooth to hop across as few as 15 1-MHz channels in the 2.4 GHz ISM band. This frequency-division approach, known as adaptive hopping, would theoretically allow modified Bluetooth devices to operate simultaneously with Wi-Fi devices by dividing the frequency band: Bluetooth would operate in one section, and Wi-Fi another, non-overlapping section. This technique is currently permissible under FCC regulations for radios operating under at or below -1.3 dBm of transmit power. The regulation must be changed, however, to allow the typical class 1, 2, and 3 Bluetooth devices to operate in this mode. This represents a significant change to the ISM band rules and requires much more explanation than allowed by the scope of this paper. However, we have provided a brief overview of several important aspects of this petition and the adaptive hopping approach.

⁶ See footnote 5.

Analysis

Adaptive hopping will provide a viable and important solution to 802.11b and Bluetooth coexistence, provided it is quickly ratified through the appropriate regulatory processes, and its recommended implementation of its intelligent adaptive hopping algorithms is well thought out. The timeliness of the regulatory process is primarily a function of the integrity of the adaptive hopping petition. If it adequately addresses the potential issues discussed in the following sections, the process should be relatively quick. However, if the petition's implementation recommendations are ambiguous, and do not adequately address the issues below, the regulatory process will likely be protracted while these important details are resolved. Regardless of the ratification timing, when passed, Bluetooth adaptive hopping will provide an excellent coexistence solution in environments with two or fewer Bluetooth piconets and no overlapping Wi-Fi networks.

Adaptive Hopping as Optional Profile (Operational Mode)

Bluetooth specification 1.0 and 1.1 do not require adaptive hopping functionality and therefore use all 79 available 1-MHz hop channels. This, of course, creates interference for any closely located Wi-Fi network, and is the core of the issue addressed by adaptive hopping. Over the next year, Cahners In-Stat estimates approximately 30 million Bluetooth devices will come to market, all theoretically under Bluetooth specification 1.1. Given that future Bluetooth devices will need to communicate with these legacy devices, the adaptive hopping petition is likely to be passed as an optional function or profile. Therefore, Bluetooth developers will not be required to include the adaptive hopping profile and may choose to eliminate it to get to market faster or optimize their product cost.

The complication arises from the fact that, when it comes to hop pattern, Bluetooth piconets must operate at the lowest common denominator. Thus, a Bluetooth device with the optional adaptive hopping functionality will be forced to bypass the mechanism and use all 79 channels if there is even one unmodified Bluetooth device in its piconet. Furthermore, since the majority of current and foreseeable Bluetooth implementations will perform the hop selection in the hardware, it will be very difficult to retroactively modify them for adaptive hopping. Accomplishing the modification would probably require a new spin or release of the Bluetooth device hardware.

This series of timing issues represents a significant consideration for adaptive hopping's universal effectiveness as a coexistence solution. There are other, more technical considerations that must also be addressed for adaptive hopping to achieve its full potential.

Adaptive Hopper Must Accurately Sense and Respond to Interferers

It appears that adaptive hopping Bluetooth devices will move into adaptive hopping mode based on one or more of several possible interference detection mechanisms. The petition does not currently specify which approach will be

recommended in the Bluetooth profile, but four of the many possible technical approaches are:

- 1) The Bluetooth device gradually adapts its normal operation hop pattern based on observed packet loss;
- 2) The Bluetooth device detects and assesses received signal strength across its wireless environment before commencing operation;
- 3) The Bluetooth device transmits a “test” pattern of packets across the entire spectrum, observes the ratio of lost packets across available channels and locates its adapted piconet in the least active or interference-prone channel; and
- 4) The Bluetooth device is collocated with a Wi-Fi device, and can receive the Wi-Fi passband location from the Wi-Fi device, so it simply avoids operating within the Wi-Fi passband.

Given that under the current petition the adaptive hopping Bluetooth device must reassess its restricted mode every 30 seconds, all of these approaches will eventually result in low interference operating scenarios. However, with the exception of the last, they all have the potential to exacerbate interference problems under certain conditions.

Bluetooth™ Difficulty in Detecting Wi-Fi™ Signal

If the Bluetooth device is attempting to modify its normal 79 MHz hop pattern based on observed packet failure rates, it could create significant interference for a closely located Wi-Fi device, but have difficulty detecting interference to its own signal, and therefore not move into adaptive mode. This is caused by a combination of factors, including Wi-Fi transmit power, the distance between the Bluetooth and Wi-Fi devices, and the Bluetooth receive filter.

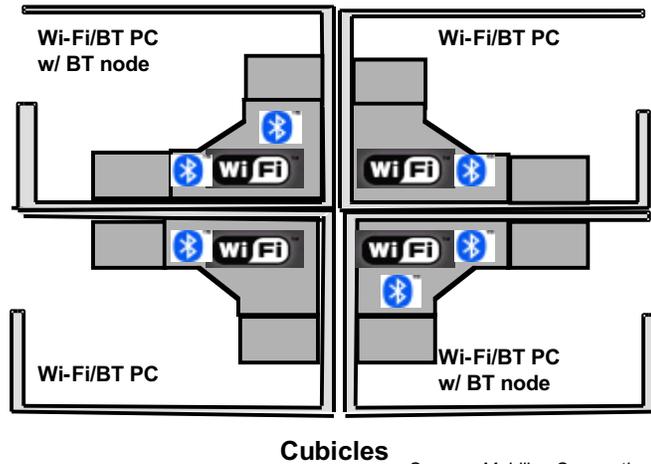
Consider the following scenario using the path loss model from A. Kamerman’s 1999 work⁷. A Wi-Fi AP and station are transmitting at +15 dBm and are separated by approximately 15 meters, resulting in received signal strength of – 52.5 dBm at either device. If a 0 dBm Bluetooth piconet is within 5 meters of either Wi-Fi device, it will create debilitating interference for the Wi-Fi network but will not sense interference to its own signal until it is within approximately 1.3 meters of the Wi-Fi transmitter (AP or station).

This occurs because, with 5 meters separation between the Bluetooth and Wi-Fi device, the Bluetooth signal reaches the Wi-Fi device at approximately – 54 dBm, resulting in an intolerable signal-to-noise ratio. On the other hand, the Bluetooth receive filter’s noise reduction effects on the Wi-Fi signal allow the Bluetooth piconet to continue operating without significant error rates, even when it is within ~ 1.3 meters of the +15 dBm Wi-Fi transmitter. The Bluetooth receive filter reduces the Wi-Fi signal by 13 dB; therefore, at 1 meter from the Wi-Fi transmitter, the +15 dBm transmission is perceived by the Bluetooth as – 38 dBm⁸. At ~ 1.3 meters, the perceived Wi-Fi signal (noise) reaches approximately – 40 dBm, and the Bluetooth device suffers increased packet error rates and can modify its hop pattern. This creates significant opportunities for continued interference, particularly in wireless corporate scenarios with multiple networks, as depicted in Figure 5.

⁷ Please see Appendix 2 for explanation of the path loss models employed.

⁸ + 15 dBm – 40 dB – 13 dB = – 38 dBm

While multiple network environments are not commonplace in today's market, Figure 5 shows a very likely corporate environment in which multiple Wi-Fi networks and Bluetooth piconets will likely arise⁹. In this scenario, enterprise cubicles are configured back-to-back with Wi-Fi/Bluetooth-enabled PCs in each corner, and with each user having an active or potentially active Bluetooth piconet with their PDA/cell phone or Bluetooth enabled peripherals.



Source: Mobilian Corporation

Figure 5 – Basic Geometry of Bluetooth and Wi-Fi Penetrated Corporation

Congested Wireless Environments are Particularly Troublesome

In a congested wireless environment, a newly introduced adaptive hopping Bluetooth device could have difficulty identifying the optimal location for its piconet. This is particularly relevant in situations where an existing adapted Bluetooth piconet already exists.

As mentioned earlier, interference to Bluetooth is greatest when the collocated Bluetooth radio is receiving while the Wi-Fi is transmitting, although this is not usually the case. In a typical wireless network usage scenario, a Wi-Fi station, such as a desktop or laptop PC, receives far more information from the access point (AP) than it transmits. This is known as an asymmetric usage model and is widely accepted as the dominant usage scenario for many networks.

Network asymmetry lessens the likelihood of interference from Wi-Fi to the Bluetooth because it is more likely that the Bluetooth is being subjected to a non-interfering, attenuated signal from the Wi-Fi AP rather than a relatively strong signal from the closely located Wi-Fi station. The Wi-Fi station transmits infrequently, and generally in the form of relatively short ACKs (less than 0.2 milli-seconds). This generates an interesting dichotomy wherein the adaptive hopper is likely to “choose” to operate directly in a Wi-Fi passband if there is already a BT adaptive hopping network in the environment. We explore this

⁹ Figure 5 shows a potential corporate deployment/configuration of Wi-Fi and Bluetooth, developed in collaboration with a leading PC OEM. The Wi-Fi stations communicate with an Access Point (AP). The Wi-Fi stations are also equipped with Bluetooth radios communicating with another Bluetooth node (potentially a headset, PDA, etc.) Additional Bluetooth piconets can be active in adjacent cubicles.

Adjacent-Channel Noise ¹⁰

Adaptive hopping does not adequately address adjacent channel, or out-of-band, interference. Adjacent channel noise is primarily related to the degree of isolation achieved between the two radios' signals. Once again, a Bluetooth device in transmit mode significantly degrades the reception of a collocated Wi-Fi receiver. This is the case whether the Bluetooth is transmitting in the same channel the Wi-Fi is receiving in or in an adjacent one. The overall effects of adjacent channel interference become more severe as the Wi-Fi station is moved further from the transmitting Wi-Fi AP, and the subsequent signal-to-noise ratio becomes smaller, rendering the attenuated AP signal more susceptible to interference.

Number of Channels

Without careful consideration of the number of channels selected as the “adaptive hopping mode,” and how and where they are spaced throughout the band, Bluetooth could continue to create problems. The 15 channels requested in the petition could result in Bluetooth being concentrated in an extremely small portion of the 83.5 MHz-wide 2.4 GHz ISM band. If this were the case, due to the resulting concentrated Bluetooth traffic, Wi-Fi communications in any overlapping channels would become nearly impossible. Also, according to the Simon, Omura [MSJO85] frequency capacity model, Bluetooth would be limited to a single piconet within any given contiguous 15-MHz channel, versus the 8-10 piconet capacity under current regulations.

MAC-level Switching

Overview

MAC-level switching describes switching functionality at the baseband level. The solution is either integrated into the two protocols' basebands or in a self-contained module that communicates with both basebands and provides switching functionality “remotely.” Mobilian presented this “remote” MAC-level functionality as a proposed coexistence mechanism at the November 2000 meeting of the IEEE 802.15.2 task group. The proposal, dubbed MEHTA, was very well received and will be voted on as a recommended best practice in the coming months.

Analysis

MAC-level switching is performed in the baseband and basically performs the same functionality as driver-level switching, but at a much faster rate and with predictable latency. Consequently, it is able to mitigate many of the interference factors that driver-level switching cannot. MAC-level switching does not suffer

¹⁰ For further explanation of adjacent-channel noise, and in-band and out-of-band noise, please see the appendix.

from transmitting signals into incoming receptions, Bluetooth polling, or operating system latency. However, like many of the previous approaches, it is susceptible to adjacent-channel interference and therefore does suffer some degradation. Also, a MAC-level approach will have a long development cycle time relative to a driver-level switching approach, and is pertinent only in systems employing both technologies in a very small, if not integrated, area. Such systems also typically have long development cycles.

Simultaneous Operation

Overview

Simultaneous operation is the ability for different, fully standards-compliant wireless systems to operate simultaneously in a collocated scenario while experiencing minimal or no performance degradation. Simultaneous operation also refers to devices able to offer the user outstanding performance without a list of operational caveats. The device should “just work,” regardless of other devices within its operating environment.

Simultaneous operation of Wi-Fi and Bluetooth will occur more and more frequently as users begin completing everyday tasks such as copying or printing a file from their Wi-Fi PC while using a Bluetooth-enabled mouse and keyboard. Its frequency will continue to grow as personal communication devices and synchronization activities with PCs and networks grow, and it will gain even more importance as distributed applications take off – “the next big thing in software” – and Bluetooth devices must coexist with massive amounts of Wi-Fi network activity.

In all these scenarios, users will appreciate being able to use whatever wireless devices surround them, when they want to and how they want to. Users will demand “Coexistence without Compromise”™, and will resist adopting wireless devices as long as there are operational difficulties or perceived concerns.

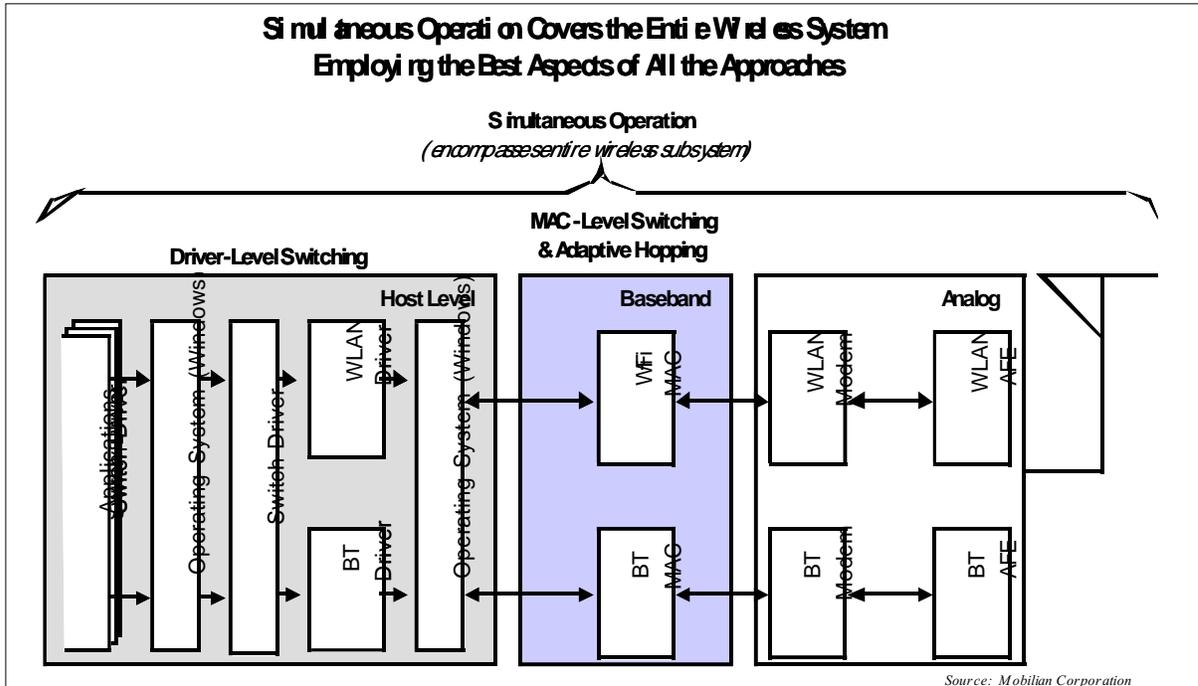


Figure 7 – Simultaneous Operation Covers Entire Conceptual Wireless System Diagram

We believe that true, sustainable simultaneous operation can only be achieved by taking a system-level approach encompassing the entire wireless sub-system. This design methodology allows the solution to selectively use the best aspects of each technique or techniques, depending on its environment and the required usage models.

For example, a driver-level switching technique may generate the best user experience in a low bandwidth synchronization scenario, while MAC-level switching will manage interference much more effectively for SCO traffic, or when a user has wireless peripherals such as speakers or a keyboard. Further, in future environments of distributed computing / applications, a system-level approach may be required to effectively allow SCO traffic while manipulating server-side office productivity applications.

Mobilian Corporation developed its first product, TrueRadio™, using a system-level methodology, and will continue development under this methodology with subsequent, multi-standard radio products.

Analysis

Collocated Mobilian™ TrueRadio™ Performance Ganymede Chariot (NetIQ)

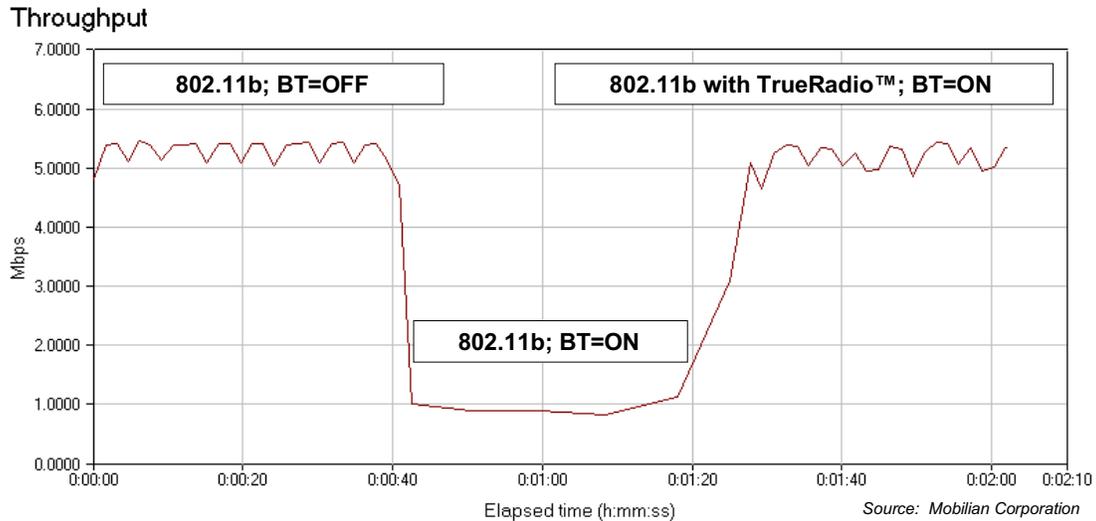


Figure 8 – Ganymede Chariot Graph of Mobilian Corporation’s TrueRadio™ Demonstration

Mobilian Corporation’s TrueRadio™ technology allows simultaneous operation by using technical enhancements across every aspect of the wireless sub-system, from antenna to application. It incorporates many of the best characteristics of the techniques described above, as well as other advanced technology not addressed here. By doing so, it is able to provide fully-standards-compliant Wi-Fi and Bluetooth radios capable of seamless, transparent operation under virtually any scenario, all without compromise to the end-user. Additionally, because of the high degree of integration in the Mobilian™ TrueRadio™ solution, the expected cost of the solution is substantially less than simple collocation of two chipsets or independent cards. Mobilian demonstrated¹¹ its TrueRadio™ technology under NDA at Comdex 2000 with very successful results and excellent market response. The TrueRadio™ technology eliminated virtually all interference as shown by the Ganymede¹² chart in Figure 8.

Summary

The market is rapidly moving toward resolving the coexistence concerns surrounding Wi-Fi and Bluetooth. The variety of approaches discussed in this paper will likely address the issue prior to it ever affecting the end-user.

¹¹ The technology demonstration used a notebook computer with collocated Bluetooth and Wi-Fi radios. The Wi-Fi AP antenna was covered in attenuating material simulating an office environment with forty-five feet between the STA and AP with 8 cubicle walls. The partner Bluetooth node was located 1 meter away.

¹² Ganymede Chariot software is used by WECA in the Wi-Fi certification process. It measures the throughput for wireless LAN systems by monitoring and calculating the time required to transmit a 1MB file from the AP to the STA. Ganymede was recently acquired by NetIQ.

Consequently, market forecasts for Bluetooth and Wi-Fi will remain strong, and a new market for combination 802.11b and Bluetooth solutions will arise. However, the need for effective, multi-standard, coexistence solutions will only increase as wireless devices proliferate and simultaneous operation usage models become pervasive. This will occur with proven, perpetual market innovation such as new power saving techniques allowing Wi-Fi to penetrate the hand-held market, and ramp of the application service provider market (distributing computing models like the Microsoft “.net” initiative). As the markets evolve, partial coexistence solutions with operational caveats and marginal user experiences will falter, and the market will demand seamless connectivity that “just works.” The Mobilian™ TrueRadio™ solution is the only end-to-end solution capable of providing users with this experience: Coexistence without Compromise™. As shown in the simulation graph below and in the Ganymede chart of Mobilian’s Comdex demonstration, the TrueRadio™ solution allows near perfect simultaneous operation of both Wi-Fi and Bluetooth even when collocated in a single card with a single antenna.

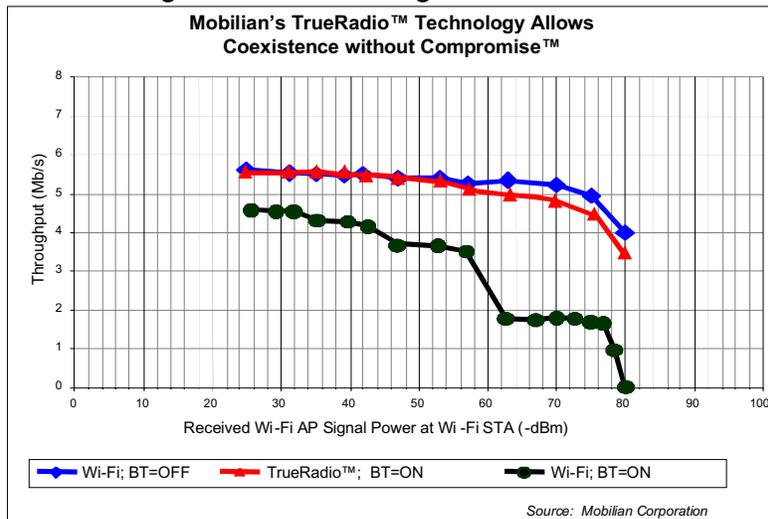


Figure 9 – Mobilian’s TrueRadio™ Performance in Collocated Scenario

Appendix 1 – In-band versus Out-of-band Noise

Signals and Noise

Every wireless communication system, by definition, consists of at least two nodes. At any given time, one node transmits (a transmitter) and the other receives (a receiver)¹³. Successful system operation depends on the receiver's ability to separate a desired signal from an undesired signal. This depends on the ratio between the energy of desired signal and the total noise (interference) at the receiver's antenna. This ratio is referred to as E_b/N_t (energy per bit over total noise) or SNR (signal-to-noise ratio). The receiver's job is to maximize the ability to decode desired signals while minimizing the ability to allow undesired signals (noise) to interfere. One of the most important characteristics of a communication system is the minimum SNR at which the receiver can still successfully decode the signal (the E_b/N_t threshold of the system). The lower the E_b/N_t threshold, the greater the system's immunity to interference. The lower the SNR, the more likely the undesired signal will cause unacceptable errors in data packets which force retransmission (and delays inherent in that process), or impact voice quality. There are also situations where noise is so strong that the receiver cannot begin to recover the desired signal.

Types of Noise

The noise at the receiver's antenna can be divided into two categories defined below and illustrated in Figure 10.

Out-of-band noise – undesired energy in frequencies that the transmitter does not use; and

In-band noise – undesired energy in frequencies that the transmitter used to transmit the desired signal

¹³ Both 802.11 and Bluetooth stations transmit or receive (half duplex systems); there are other systems where a station transmits and receives at the same time (full duplex systems); the discussion below applies to both type of systems.

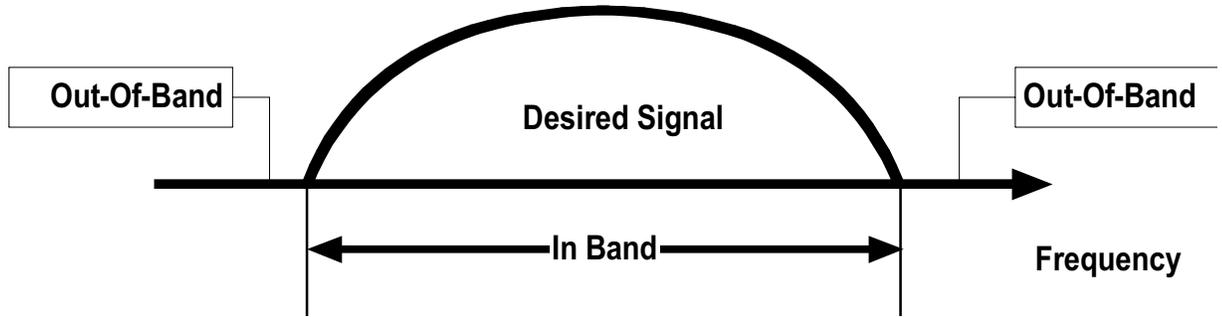


Figure 10 – In-Band versus Out-Of-Band Noise

Both in-band and out-of-band noise can degrade a wireless communications system’s performance. Out-of-band noise can usually be filtered out because the energy in the system’s frequency band does not carry any useful information. In-band noise, is much more problematic.

Noise can be further categorized as either “white” or “colored.” White noise is a collection of energies transmitted from many different sources without any coordination between them. This energy is typically distributed evenly across the frequency band and does not have any deterministic behavior over time or frequency. Colored noise is transmitted from intentional radiators and has a specific behavior in time and frequency. Figure 11 * MERGEFORMAT illustrates the difference between white and colored noise.

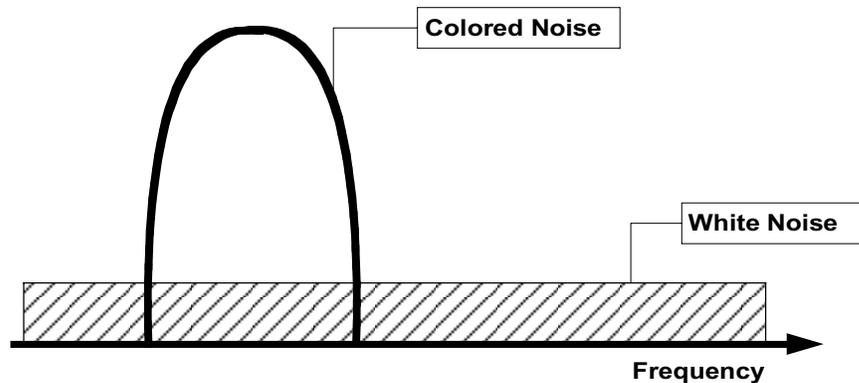


Figure 11 – White Noise and Colored Noise are Very Different

Most wireless communication systems assume that the only type of in-band noise is white noise. Other intentional radiators are assumed to transmit out-of-band. Receiver designs with their associated filtering techniques are optimized around these assumptions. Unfortunately, in a case where two intentional radiators such as Bluetooth and Wi-Fi both share the same frequency band, receivers must also address the case of in-band, colored noise.

Every transmitter is supposed to transmit only within a limited bandwidth; however, this is not physically possible without injecting noise into adjacent frequencies (sideband signals), as shown in Figure 12. The amount and nature

of sideband signals created during transmission are determined by what is referred to as the transmitters' "transmit mask." Sideband signals must be considered when evaluating interference between wireless systems sharing the same frequency band. In addition, receiver filters cannot be perfectly "rectangular." This means that the filter cannot precisely differentiate between signal and noise just inside and outside the passband. The combined impacts of transmit and "receive" masks, explain what is referred to as adjacent channel interference.

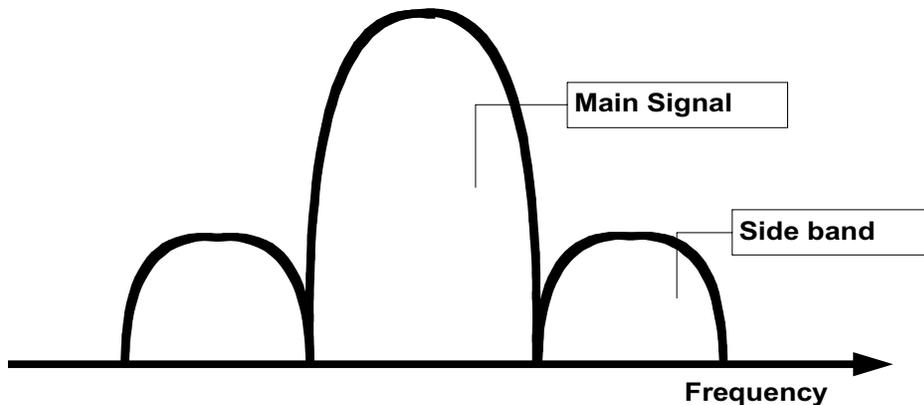


Figure 12 – Typical Transmit Mask

Bluetooth and Wi-Fi Interference Cases

Bluetooth and Wi-Fi share the same band. **If Bluetooth and Wi-Fi operate at the same time in the same place, they will interfere with each other (collide).** Specifically, both systems transmit on overlapping frequencies¹⁴, creating in-band colored noise for one another. Interference between Bluetooth and 802.11b occurs when either: a) an 802.11b receiver senses a Bluetooth signal at the same time as an 802.11b signal is being sent to it, so that the effect is most pronounced when the Bluetooth signal is within the 22 MHz-wide passband of the 802.11b receiver; and b) a Bluetooth receiver senses an 802.11b signal at the same time as a Bluetooth signal is being sent to it, so that the effect is most pronounced when the 802.11b signal is within the passband of the Bluetooth receiver. These time-frequency collisions are summarized in the following table:

¹⁴ We also need to take into account the sidebands of each transmission.

		Bluetooth Tx		Bluetooth Rx	
		In-band	Out-of-band	In-band	Out-of-band
Wi-Fi™	Tx	No Conflict	No Conflict	Strong Interference	Moderate Interference
802.11b	Rx	Strong Interference	Moderate Interference	Strong ¹⁵ Interference	Moderate Interference

Source: Mobilian Corporation

Table 3: The Interference Cases for Bluetooth and Wi-Fi¹⁵

As we discussed earlier, the impact of interference varies as a function of energy ratios. Therefore, the “strong” versus “moderate” interference representations in Table 3 should be viewed as “relative” effects. In general, in-band interference impacts are more deleterious than those of out-of-band. However, the impact of both classes of interference depends upon SNRs, which vary according to a number of parameters discussed in more detail later.

It is worthwhile to note that neither Bluetooth, nor Wi-Fi was designed with specific mechanisms to combat the interference they create for each other. As a fast, frequency-hopping system, Bluetooth assumes that it will hop away from bad channels, minimizing its exposure to interference. The 802.11b Media Access Control (MAC) layer, based on the Ethernet protocol, assumes that many stations share the same medium, and if a transmission fails it is because two Wi-Fi stations tried to transmit at the same time. Later, we will discuss how this assumption drives system behavior that actually worsens the impact of interference from Bluetooth.

¹⁵ See Footnote 5.

Appendix 2 – Path Loss Models Employed

The idea of proximity is critical for this analysis. Since radio-frequency signals expand outward from the antenna in roughly spherical waves, the signal strength decreases rapidly with the distance from the transmitting antenna. This path-loss profile is proportional to the reciprocal of the squared distance from the antenna at short ranges, and attenuates more sharply in indoor environments at a distance of a few meters from the antenna, as reflection effects become apparent. The following path-loss models (L_p in dB) will be used [Kam99]:

$$L_p = -40.0 - 20\log(d), \text{ when } d \leq 8 \text{ meters}$$

$$L_p = -58.5 - 33\log(d/8), \text{ when } d > 8 \text{ meters}$$

Using these empirical models, the loss profile can be plotted vs. distance as shown in Figure 13. A communication system must overcome this path loss using a combination of transmission power and receiver sensitivity to meet a desired signal-to-noise or signal-to-interference ratio criterion, which in turn translates into a probability of error. If the received signal strength is not adequate to overcome either noise or interference, unacceptable packet errors result and system throughput and latency will deteriorate.

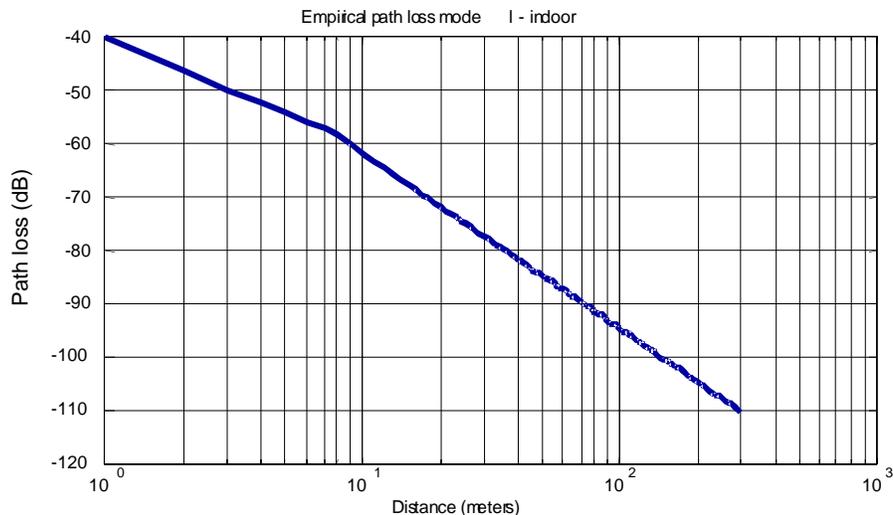


Figure 13 – Path Loss as a Function of Distance, Indoor, 2.4-GHz ISM Band

References

[CIS00a] Cahners In-Stat, *Bluetooth 2000 – To Enable the Star Trek Generation*, July 2000

[CIS00b] Cahners In-Stat, based on preliminary data to be published in December 2000

[MBLN01] Mobilian Corporation, *Wi-Fi™ (802.11b) and Bluetooth Simultaneous Operation: Characterizing the Problem*, November 2000

[NGFM00] Nada Golmie and F. Mouveraux, *WPAN Coexistence Performance Evaluation: MAC Simulation Environment and Preliminary Results*. IEEE 802.15-00/066r0, March 2000.

- [SS00] S. Shellhammer, *Packet Error Rate of an IEEE 802.11 WLAN in the Presence of Bluetooth*. IEEE 802.15-00/133r0, May 2000.
- [GE98] G. Ennis, *Impact of Bluetooth on 802.11 Direct Sequence*. IEEE 802.11-98/319, September 1998.
- [JZ98] J. Zyren, *Extension of Bluetooth and 802.11 Direct Sequence Model*. IEEE 802.11-98/378, November 1998.
- [JZ99a] J. Zyren, *Reliability of IEEE 802.11 Hi Rate DSSS WLANs in a High Density Bluetooth Environment*. Bluetooth '99, June 8, 1999.
- [JZ99b] J. Zyren, *Reliability of IEEE 802.11 WLANs in Presence of Bluetooth Radios*. IEEE 802.15-99/073r0, September 1999.
- [BT99] Specification of the Bluetooth System, version 1.0B, December 1, 1999.
- [Kam99] A. Kamerman, "Coexistence between Bluetooth and IEEE 802.11 CCK Solutions to Avoid Mutual Interference" Lucent Technologies Bell Laboratories, Jan. 1999, also available as IEEE 802.11-00/162, July 2000.
- [MSJO85] M.K. Simon, J. Omura, R. Scholtz, and K. Levitt, *Spread Spectrum Communications Vol. I, II, III*. Rockville, MD. Computer Science Press, 1985.
- [ETSI99] ETR0230002 V0.2.0 (1999-04) Broadband Radio Access Networks (BRAN); High Performance Radio Local Area Networks (HIPERLAN) Type 2; System Overview (www.etsi.org)
- [FCC15.247] Code of Federal Regulations, Title 47, Volume 1, Part 15, Subpart C (Intentional Radiators), Section 247. US Government Printing Office. October 1, 1998 Revision of 47CFR15.247, pages 690-692. Available on the web at: http://www.access.gpo.gov/nara/waisidx_99/47cfr15_99.html.