Reconsidering Wireless Systems with Multiple Radios

Paramvir Bahl, Atul Adya, Jitendra Padhye, Alec Wolman
Microsoft Research, 1 Microsoft Way, Redmond, WA 98052

ABSTRACT

The tremendous popularity of wireless systems in recent years has led to the commoditization of RF transceivers (radios) whose prices have fallen dramatically. The lower cost allows us to consider using two or more radios in the same device. Given this, we argue that wireless systems that use multiple radios in a collaborative manner dramatically improve system performance and functionality over the traditional single radio wireless systems that are popular today. In this context, we revisit some standard problems in wireless networking, including energy management, capacity enhancement, mobility management, channel failure recovery, and last-hop packet scheduling. We show that a systems approach can alleviate many of the performance and robustness issues prevalent in current wireless LAN systems. We explore the implications of the multi-radio approach on software and hardware design, as well as on algorithmic and protocol research issues. We identify three key design guidelines for constructing multi-radio systems and present results from two systems that we have built. Our experience supports our position that a multi-radio platform offers significant benefits for wireless systems.[1]

1. INTRODUCTION

Wireless communication has become an indispensable part of modern day-to-day life. As a consequence, the importance of building robust wireless systems cannot be overstated. Unfortunately, current wireless communication systems are not robust. Ask anyone who uses a wireless network regularly, whether in a corporate setting or in the home, and they will relate experiences of dropped connections; problems with authentication; reduction of throughput due to interference from other wireless devices; and loss of use of a device due to rapid battery drainage.

Previous attempts to rectify this situation have met with limited success. A fundamental shortcoming of many previous attempts is the implicit assumption that the underlying system has only a single radio, or in a few cases multiple radios that operate independently.

We believe that future solutions to achieving robust wireless systems do not have to rely on requiring major breakthroughs in radio technology, but instead require careful thinking about the demands of the task to be accomplished, an understanding of the limitations of the available components and the interactions between them, and innovation in combining these components in a way that uses their strengths constructively. The problems cannot be solved at the level of an individual component, rather they must be solved by considering the behavior of the complete system. Thus, the challenge is to design systems that use all available resources (spectrum, energy, hardware, and software) efficiently while being general enough to support a wide variety of applications.

[1]This paper is a product of over 3 years of research on multi-radio wireless LAN and wireless mesh systems.

It is our thesis that systems using multiple radios on each node, working in an integrated manner to accomplish a common task, provide significant benefits in terms of functionality and performance over systems with a single radio, or with multiple radios where the radios are used independently.

A current trend in the IEEE 802.11 [12] wireless LAN industry is to build products that support multiple standards. For example, both Netgear and Orinoco sell client adapters that support IEEE 802.11{a, b, and g}. Although these products contain multiple radios, they do not meet our definition because the multiple radios cannot be used in an integrated manner. For example, when the card is operating in IEEE 802.11a mode, it cannot maintain connectivity to an IEEE 802.11b or an IEEE 802.11g network simultaneously.

Today's wireless LANs are built using radios that support high data rates, but these radios consume a large amount of energy [8, 25]. Wireless personal area networks (PANs) use radios that are energy efficient, but support low data rates [9, 17]. Systems containing both types of radios are starting to emerge in the market. However, networking stacks treat these radios as two distinct network interfaces and broad classes of applications use these interfaces independently. We show that a system combining both radios in an integrated manner can provide the benefits of both – high data rates with low energy consumption.

Our goal is to show that many standard problems in wireless networking can be solved in a much better way if the wireless platform includes multiple radios and multiple physical channels. Not only are multiple-radio systems desirable in terms of the functionality they offer, they are also practical in terms of the component costs, because of impressive decreases in the cost of individual radios.

In advocating the use of multi-radio systems, we do not recommend any particular radio technology. In some cases, it is desirable to have radios that differ in their physical characteristics, while in others multiple similar radios are useful. By adding appropriate support to the operating system, unmodified applications can take advantage of the increased functionality provided by such a system. On an as needed basis, the software either chooses the most appropriate radio or uses multiple radios simultaneously, depending on the task that needs to be accomplished.

In the remainder of this paper, we investigate multi-radio solutions to five standard problems in wireless networking, and present our experiences building such systems for two of these problems. One system uses multiple radios to provide better energy management in a wireless LAN system; another uses multiple radios to improve capacity in multihop wireless mesh system. We also outline multi-radio solutions for mobility management, channel failure recovery, and last-hop packet scheduling. Our objective is to make a compelling case that a systems approach based on an integrated multi-radio platform can improve both functionality and performance of future wireless systems.

2. EXISTING WIRELESS TECHNOLOGIES

What follows is a very brief summary of existing wireless LAN and PAN technologies. The IEEE 802.11 standard, commonly referred to as the "WiFi Standard", dominates the wireless LAN industry world-wide. There are three popular variants of this IEEE standard: 802.11a, 802.11b, and 802.11g. IEEE 802.11a operates in the 5 GHz unlicensed frequency band whereas 802.11b and 802.11g operate in the 2.4 GHz unlicensed frequency band. 802.11a and 802.11g support a maximum data rate of 54 Mbps, whereas 802.11b supports a maximum data rate of 11 Mbps. The three standards differ in their choice of signal modulation but are similar in their choice of carrier sense multiple access technology as their medium access control (MAC) protocol. In Europe, there is another wireless LAN technology called Hi Performance LAN/2 (HIPERLAN/2) that provides up to 54 Mbps data rates in the 5 GHz band. Radio modulation is similar to the IEEE 802.11a & g but the MAC is based on dynamic time division multiple access.

In the wireless PAN space, the Bluetooth standard, also operating at 2.4GHz, dominates. The IEEE recently standardized a low-power, low data rate radio as the IEEE 802.15.4 (Zigbee) standard [9]. Zigbee operates in two different frequency bands including the 915 MHz and the 2.4 GHz bands and supports comparatively low data rates of up to 220 Kbps. In addition to the standards, there are many other niche market radios available today, such as the RF Monolithics TR1100 radio [23], a low-power radio which operates at 915MHz and provides data rates up to 1 Mbps.

The main point of this review is that there is no "one shoe fits all" radio technology. There are a wide variety of radio technologies available to us today, each focused on certain classes of applications, each making certain tradeoffs. We propose that system designers combine these technologies in a way so that the strengths of one radio can overcome the weakness of other radios to improve the overall performance of the system.

3. THESIS AND DESIGN GUIDELINES

Our thesis is: Wireless systems performance and functionality can be improved significantly by designing a system that employs multiple radios in the same network. Although these radios provide different communications links, the networking protocols and algorithms utilize the radios in conjunction to accomplish the same task. Such a system will be more dependable, more flexible, and allow more innovation than traditional single radio wireless systems.

Within this framework, we identify three specific design guidelines:

- **Design for Choice**. Different radios with different properties and characteristics should operate inside the same platform. The diversity yields significant system performance advantage.
- **Design for Flexibility**. Programming abstraction should allow applications to treat the radios as a single logical pipe, yet allowing networking protocols to access the radios directly. The flexibility encourages innovation at all levels of the networking stack.
- **Design for Separation**. Where possible, separate the control and the data traffic. Separate radios may be dedicated to each of the functions.

It is important to realize that additional hardware does not always improve system performance. When resources are shared, more hardware, without proper arbitration and safeguards can lead to lower performance. Wireless systems provide a good example of this: the spectrum or the communications link is a shared resource. Consequently, there is a limit beyond which more hardware will not help [1].

4. REVISITING STANDARD PROBLEMS

We discuss five important problems in wireless networking: energy consumption, network capacity, mobility management, communications robustness, and quality of service. We show that by using a multi-radio approach, system designers can deliver better solutions to these classic problems.

Note, multiple radio systems where the radios work independently are available today. Examples include the Atheros AR5001X Dual-Band 802.11a/b/g Wireless LAN Network Card [5] and the HP iPAQ h6315 Pocket PC - Phone Edition [18], which contains a Bluetooth radio, a WiFi radio and a quad-band GSM/GPRS radio. In such systems, multiple radios might even work simultaneously, but there is usually no explicit co-ordination of their activities. The systems we propose in this paper are designed explicitly to co-ordinate the activities of multiple radios, to optimize overall system performance.

In the discussions that follow, we make two important assumptions: First, we assume that spectrum will always be limited and as a consequence, from Shannon's theorem [24], channel capacity will also be limited. Second, we assume that the cost of including additional radios in a wireless system will continue to decline dramatically. Given the current trends in wireless technologies, both these assumptions are reasonable.

Finally, the ideas explored in this paper open new areas of research in protocol design, algorithms, and systems development. It is our hope, that others will share our enthusiasm by continuing research in these areas, which in turn will help deploy robust multi-radio systems in everyday life.

4.1 Energy Consumption

Several previous studies have established that high data rate wireless LAN systems are energy inefficient (see [8] and the references therein). Consequently, researchers have proposed different techniques to improve energy efficiency. The various approaches to the problem can be broadly classified into the following categories: improving channel access mechanisms [26]; maximizing sleep times of wireless cards in accordance with the traffic pattern [14]; and using the lowest possible power level for transmitting and receiving data [3]. There is also on-going research looking for ways to reduce the energy consumption of different RF circuitry components [21].

All of the prior work has focused on the single radio case and the effciency improvements have not been as dramatic as we would like them to be. Consequently, energy inefficiency continues to limit the use of wireless LAN technology in battery powered mobile devices. For example, a PDA-based WiFi phone is impractical since it drains the battery at a rapid rate even in "standby mode". Furthermore, battery capacities are unlikely to improve significantly in the near future [20].

With current 802.11 wireless technology, a significant amount of energy loss occurs while the network card is in idle mode, i.e., while it is waiting for messages to arrive. Based on this observation, we propose the use of an additional low-power radio (LPR) along with a high-power 802.11 radio (HPR) for achieving energy efficiency in a wireless LAN system.

Figure 1 shows the idle power consumption of two commercially available radios – the TR1000 low-power radio [23] and a Cisco AIR-PCM350 802.11b radio. The TR1000 consumes significantly lower power than the 802.11b radio. Observing this, we briefly describe three approaches that we have developed for using an additional low-power radio in a wireless LAN system. These approaches adhere closely to the three design principals we discussed in Section 3.

A note of caution, radios using diverse communication technologies may have different ranges. When employed together the mismatch in radio ranges can lead to situations where the performance of a multi-radio system is no longer compelling. We are exploring techniques to handle such scenarios. In the meantime, we note that in the case where all radios on the sender can communicate with all radios on the receiver, our system performs significantly better than a single radio system. In the case where range mismatch prevents all radios on the sender to communicate with the radios on the receiver, our system reverts to a single radio mode.

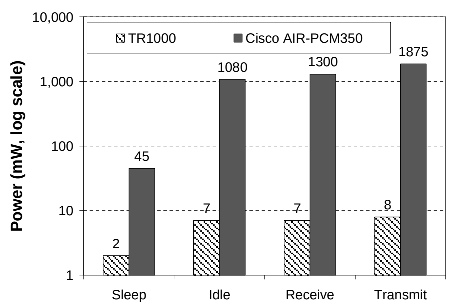

Figure 1: ***Power consumption by the TR1000 low-power radio and an 802.11b card***

4.1.1 Alert-On-LPR

Our first approach is called *Alert-on-LPR*. In this we use the low-power radio to wake up the mobile device when data is available for it. Since we previously published a detailed description of this approach [25], we provide only a short summary here.

When the mobile device is not in use, it is shut down along with the high-power wireless network card; the low-power radio is kept on. When some data is available for the mobile device, a *wakeup* message is sent to the LPR. The LPR wakes up the device along with the high-power network card and the rest of the protocols are executed using the HPR. Thus, instead of spending large amounts of energy on the HPR waiting for a message, the client expends lower energy by idling on the low-power radio. The Alert-on-LPR strategy is useful for alert-based applications such as a PDA-based phone which can spend a significant fraction of the time waiting for calls.

Figure 2: A multi-radio system we built (a) Compaq IPAQ with a WiFi radio and a ASH TR1000 transreceiver (b) A ASH TR1000 transreceiver that plugs into the serial port of any computer

Proof of Concept

We implemented the Alert-On-LPR strategy by building a WiFi phone into a commercially available PDA (see Figure 2). Our phone shuts down the entire device except the LPR circuitry, using it to listens on the LPR for any incomping calls. The client part of the implementation consists of the TR1000 low-power radio (915 MHz frequency band) along with micro-controller that controls the radio and for sending/receiving messages when the PDA and the 802.11b card (high-power radio in the 2.4 GHz band) are off. Before powering down, the device registers itself with an *LPR proxy* that can communicate with the mobile device using the TR1000 radio. The LPR proxy sends a registration request to a presence server and registers itself on behalf of the device. When a lookup for the device is received by the presence server, it sends a request to the LPR proxy who wakes up the PDA by communicating on the TR1000 radio. At this point, the device starts using the 802.11b card for normal communication.

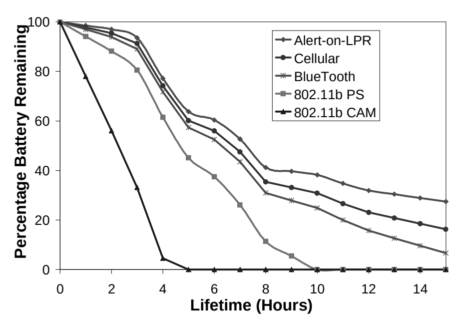

Figure 3: ***Battery lifetime of a PDA-based phone with real cell-phone usage data***

Our experiments show that using the Alert-on-LPR strategy increases the "standby time" (i.e., duration of time that the battery lasts while waiting for calls) of a PDA-based phone significantly. A PDA with a regular 802.11b card in active mode (CAM) has a battery lifetime of less than 4 hours; with power-save mode (PSP), the lifetime increases to 14.5 hours. With the Alert-on-LPR strategy, the battery lifetime is more than 30 hours, i.e., much closer to the regular battery lifetime (35 hours) of the PDA without any wireless connectivity. Thus, the Alert-on-Strategy enhances battery lifetime by more than 115% over power-save mode. From a user's perspective, the battery needs to be recharged at a much lower frequency compared to a single-radio system. Figure 3 shows the battery lifetime of the system using a realistic workload obtained from the monthly cell phone bills of two users.

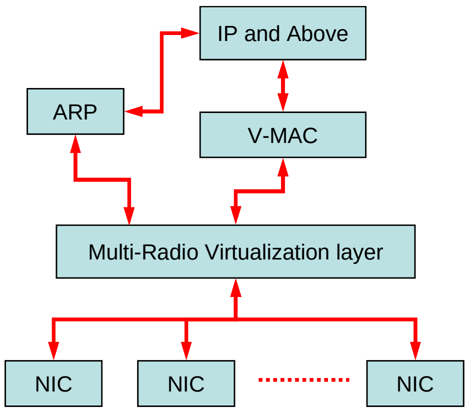

Figure 4: ***Multi-Radio Virtualization Architecture***

4.1.2 Control-On-LPR

Our second approach, called *Control-on-LPR*, generalizes the previous strategy and coordinates the wakeup of the HPR at a fine-grained level using the LPR. When a mobile node sends a message and is waiting for a reply, the HPR is switched off and the device waits for a wakeup message on the LPR (to reduce observed latency, the protocol can detect periods of high traffic and not switch off the HPR). Thus, energy consumption is significantly reduced because the mobile device idles on the LPR even while it is actively communicating. Unlike Alert-on-LPR, this mechanism can be used for all applications.

4.1.3 Data-On-LPR

Our third approach, called *Data-on-LPR*, generalizes the previous scheme further and offloads some of the data communication work from the HPR to the LPR to save energy. This technique is based on the observation that the bandwidth provided by some of the current low-power radios is quite acceptable for a number of applications, e.g., the TR1100 radio can support a raw bandwidth of 1 Mbps and the upcoming IEEE 802.15.4 technology (Zigbee) [9] supports 220 Kbps. Data-on-LPR detects the current bandwidth requirements of the mobile device and uses the LPR for sends/receives if the requirements are low. By dynamically switching between the HPR and LPR according to the device's bandwidth requirements, this strategy saves energy while ensuring that the user does not perceive a decrease in performance. Note that this strategy switches radios at a coarser-granularity than Control-on-LPR. However, it can be implemented without modifying the MAC on the HPR whereas the fine-grained switching in Control-on-LPR requires firmware modifications.

Proof of Concept

We have implemented the Data-on-LPR technique as a user-level daemon process and a kernel-level driver that resides below the networking layer but above the data-link layer. The kernel-level driver exposes the two radios as a single wireless network to the upper layer protocols and applications, as shown in Figure 4. The driver sends the packets on the radio as instructed by the daemon process. When a packet is received by the driver on one of the two radios, it simply sends the packet up to the higher layer networking protocols.

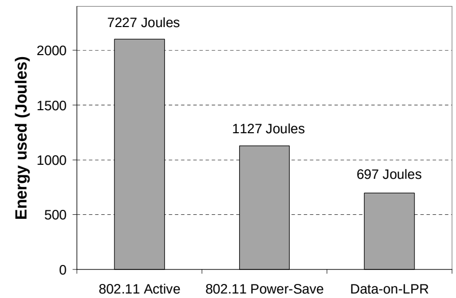

Figure 5: ***Preliminary analysis indicates that significant power savings are possible by using Data-on-LPR.***

The main purpose of the user-level daemon is to determine which radio to use and then coordinate the switching of radios. The main reason to switch from the LPR to the HPR is when an application is experiencing congestion on the LPR. Similarly, it needs to switch from the HPR to LPR when a machine is not consuming significant bandwidth and the LPR is not congested.

We have conducted some preliminary experiments to evaluate different switching strategies with an implementation and the NS-2 simulator [27]; we used network traffic collected from a developer's machine. The simulation results of our system compared with a single radio system are shown in Figure 5. The Data-on-LPR technique consumes 40% less energy than even the power-save mode (PSM) of 802.11b. By enabling PSM on the high-power radio, we believe that the energy savings can be improved further.

4.1.4 Summary

In this section we showed that by adhering to the three design principals i.e using radios with different properties and characteristics (*Design for Choice*); separating control from data (*Design for Separation*) and abstracting multiple radios to look like a single wireless network (*Design for Flexibility*), we are able to reduce the energy consumption and enhance battery lifetimes on mobile devices substantially.

4.2 Capacity Enhancement

In this section, we focus on improving system throughput in multi-hop wireless mesh systems by co-ordinated use of multiple radios. Most existing wireless LAN standards divide the available spectrum into chunks of 20 MHz, which are called "channels". A simple and effective way to improve the capacity of such a system is to use multiple radios, each of which is tuned to different non-interfering channel. Thus, a node with two radios can send packets on two channels simultaneously. More importantly, forwarding nodes can both send and receive at the same time. Using appropriate protocols for channel selection and assignment, such a system can provide significantly greater capacity than a single radio sys-

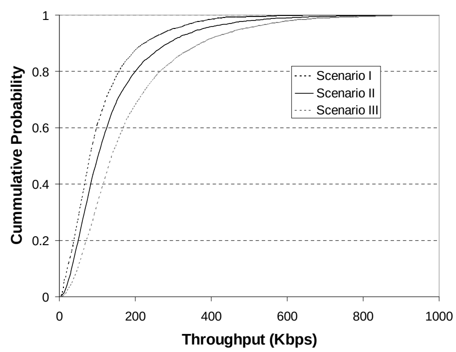

Figure 6: ***Improvement in TCP throughput. Scenario 1: every house has one radio; Scenario II: 50% of the houses have two radios; Scenario III: all houses have two radios***

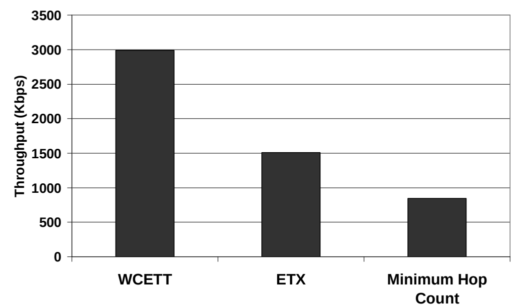

Figure 7: ***Median TCP throughput from 100 transfers in a 23-node multi-radio wireless mesh network***

tem. Later in this section, we will present results from two systems that we have built that serve as a proof of this concept.

It is not necessary that all the radios in the system operate in the same general band but different channels. Radios that work in different frequency bands may also be used, e.g., each node may have one 2.4 GHz radio and one 5 GHz radio. While co-ordinating such radios it is important to remember that the radios may have significantly different communication ranges and data rates. We are currently designing new routing protocols to fully exploit the potential of such heterogenous multi-radio systems [7].

Another way to improve capacity in multi-hop wireless meshes is to choose a combination of radios, such that one of the radios has significantly lower bandwidth, but much higher range than other radios in the system. The low-bandwidth radio on each node can then be used as a control channel to efficiently schedule data transmissions on the high-bandwidth radios. Such a system can be significantly more efficient than present systems that use contention-based MAC protocols like IEEE 802.11. We are currently conducting studies to understand the nuances involved in designing such systems.

Many researchers have explored striping across multiple network interfaces to increase throughput (see [10] and the references therein). We present two different approaches to take advantage of multiple radios per node.

4.2.1 Multi-Radio Unification Protocol

In the first approach we propose a link layer protocol called the *Multi-radio Unification Protocol* or MUP. The protocol coordinates the operation of multiple wireless network cards tuned to non overlapping frequency channels, using locally available information only.

Proof of Concept

The MUP virtualization architecture is illustrated in Figure 4. Using a prototype implementation and detailed simulations we have shown that MUP provides significant performance improvement in a wide range of scenarios (see [1] for details). Here, we present a sample result.

We simulated a a real-world topology of a suburban neighborhood consisting of 35 houses. These houses connect to each other using 2 Mbps IEEE 802.11 wireless links. Wired Internet access is available at one of the houses, and other houses share it using multi-hop routing over wireless links between the houses. We assume that there are four web users, located in four different houses, who accessed the Internet using this shared link. We simulate web traffic using the model provided in NS-2 [27]. We consider three scenarios: in the first scenario, every house has only one radio. In the second scenario, half the houses are equipped with two radios each and run MUP, and in the third scenario, all houses have two radios and run MUP.

Figure 6 shows the CDF of the throughput achieved by the four web surfers in each of the three scenarios. The CDF is calculated over all web objects. The results show that when all houses the have two radios, median throughput increases by over 70%. Moreover, even when only half the houses have two radios, the median throughput increases by almost 30%.

4.2.2 Multi-Radio Link Quality Source Routing

The MUP protocol operates using locally available information only. In the second approach, we propose a routing metric called *Weighted Cumulative Transmission Time* or WCETT, implemented in a routing protocol called *Multi-Radio Link Quality Source Routing* or MR-LQSR. MR-LQSR is a source routed, link-state protocol. The nodes in the network each have multiple wireless cards tuned to non-overlapping frequency channels. Information about bandwidth, loss rate, and the channel of every link in the network is disseminated to every node in the network. The sender uses this information to calculate the WCETT values of all available paths to the destinations, and selects the path with the minimum WCETT value. The WCETT metric is designed to favor paths that have channel diversity. Channel-diverse paths tend to provide higher throughput due to reduced interference among hops along the path. In contrast to MUP, the the WCETT metric uses *global information* to make routing decisions that are optimal on a per-flow basis. More details about WCETT and MR-LQSR are available in [7].

Proof of Concept

We have built and deployed the MR-LQSR protocol in a 23-node multi-radio, multi-hop testbed network. Our results show that our routing metric significantly outperforms previously proposed metrics such as ETX [6] and the minimum hop-count [13]. Here, we present a sample result.

In our testbed, each node was equipped with an 802.11a and an 802.11g radio. The range of the 802.11g radio is somewhat higher than the range of the 802.11a radios, and the data rate is somewhat lower. Thus, the routing metric needs to take both these factors into account. We carried out 100 2-minute long TCP transfers between randomly chosen pairs of nodes in our testbed. We repeated the test for three routing metrics: WCETT, ETX and minimum hop count. In Figure 7 we show the median throughput of the 100 transfers under each of the three metrics. We see that WCETT significantly outperforms the other two metrics.

4.3 Mobility Management

Wireless LAN standards usually do not address user mobility; it is handled by higher layer protocols. As a mobile node moves between various Access Points (APs), it first breaks the connection and then scans all the channels to find candidate APs to which it can connect. This is called an "association". During this break-scan-associate cycle, packets belonging to an on-going connection incur delay as they must wait in the mobile node's buffer. Security protocols such as IEEE 802.1x [11] further aggravate this problem because they require the client to re-authenticate with the remote security server each time it associates with a new AP. For applications such as an Internet phone or streaming video, this delay, which can be on the order of several hundred milliseconds, is not acceptable. There is vast literature on the handoff management and optimization [2, 19], but these approaches are sub-optimal as the researchers focus on single radio WLAN systems.

Ideally, if a mobile node can communicate with multiple APs at the same time, there need not be any delay or packet loss as it moves from one AP to another. Unfortunately, the physical layer of WLANs precludes such functionality.

Wireless systems with multiple radios can provide this functionality by taking advantage of the fact that in most wireless networks, APs are placed to provide some overlapping coverage. When the mobile node reaches an overlapping coverage area, the idle radios detect new APs and start associating and authenticating with them. They complete the protocol before the active radio breaks the connection with the current AP. This approach converts a "hard-handoff" (break-scan-associate-authenticate) into a "soft-handoff" (scan-preassociate-authenticate-break-associate).

With proper network engineering and protocol design, zero delay and zero packet loss can be achieved even as the mobile node moves from one AP to another. We now present a simple set of conditions that can be used for deploying access points such that no packets are lost during the reassociation/reauthentication process.

Consider a mobile with two radios, R_1 and R_2, with different ranges. Assume that R_1 is supporting an on-going connection. Suppose that the mobile is currently associated with AP_1 but is moving towards AP_2, where AP_1 and AP_2 are neighboring APs sharing some areas of overlapping coverage. Now let,

$t_b \rightarrow$ time instant at which R_2 senses AP_2
$t_h \rightarrow$ time instant at which R_2 makes the determination that AP_2 is the AP with which it should associate (AP_2 has the higher SNR value)
$t_s \rightarrow$ time instant at which R_1 breaks connection with AP_1
$t_A \rightarrow$ amount of time it takes to complete the association and authentication protocol

Figure 8 illustrates the relationship between the various events. $t_s + t_A$ is the time at which the client has resumed normal operation (note that $t_s + t_A$ is also the time when a single radio mobile would have re-established its connection after authentication).

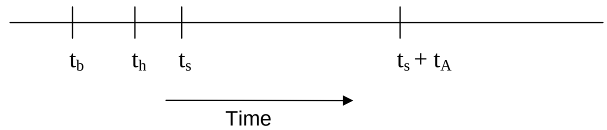

Figure 8: Time relationship between different events taking place during a handoff

To achieve zero packet loss, the client should complete the association and authentication before it breaks from old AP, i.e.,

$$t_A \leq (t_s - t_h)$$

This equation provides a value for t_A that network system engineers can design for as they consider various authentication protocols.

Now let us represent, $s \rightarrow$ mobile's speed; $d_1 \rightarrow$ distance from AP_2 up to where the first radio R_1 can stay associated with AP_1; and $d_2 \rightarrow$ distance from AP_2 at which the second radio R_2 can detect a strong signal from AP_2 (the new AP)

A desirable goal is to have the mobile detect a strong signal from the new AP (AP_2) while still being connected to the old AP (AP_1). This goal can be achieved if $d_2 > d_1$; for system designers this means that the range of the radio R_2 should be greater than the range of the radio R_1.

If t_A is determined empirically and a reasonable estimate of s is made, then the amount of overlapping coverage can be determined by using the equation $t_A < (d_2 - d_1)/s$, i.e. the time it takes the mobile to travel a distance of $(d_2 - d_1)$ is greater than the time it takes to associate and authenticate the mobile.

In summary, to ensure a handoff without packet loss, network administrators can use the following two guidelines while configuring the networks. First, the client should complete the association and authentication before it breaks from the current AP. This condition provides an upper bound on the authentication time of a security protocol. Second, the time it takes for association and authentication should be less than the time it takes for the client to travel the distance between when it first detects a strong signal from the new AP and when it breaks from the current AP. If the time for association and authentication is determined empirically and a reasonable estimate of the client's speed is made, the amount of overlapping coverage can be determined using the conditions described above.

Thus, with proper network design, a dual radio system can achieve zero packet loss and minimal packet delay which in turn is similar to what soft handoff provides. With this property, time-sensitive applications such as Internet telephony can be supported in wireless

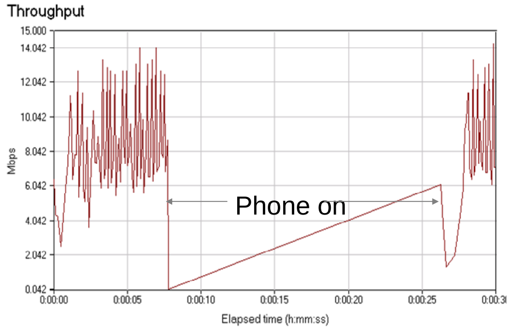

Figure 9: ***Impact on TCP throughput of a Cisco IEEE 802.11b WLAN when a 2.4 GHz wireless phone is operated close to it.***

LAN settings even as the user moves inside the network.

4.4 Channel Failure Recovery

In wireless communications there are two broad sources of channel errors. These are errors due to environmental interference and errors due to interference from other wireless networking devices. Figure 9 shows the impact on wireless LAN throughput of a nearby cordless phone operating on the same frequency band.

Graceful degradation and robustness against channel errors is possible by employing frequency diversity. Frequency diversity can be achieved by using multiple radios and operating each on different frequencies. The key insight from our approach is to build the appropriate radio switching logic in such a way that the wireless devices can employ frequency diversity transparent to the applications. Thus, a 2.4 GHz phone would not be able to destroy communication on a multi-radio wireless LAN, because whenever it causes interference in the 2.4 GHz band the software would seamlessly switch to a different band (e.g. the 915 MHz or the 5 GHz band) and continue normal operation.

4.5 Packet Scheduling

IEEE 802.11 uses a contention-based MAC. When a client wants to transmit a packet, it senses the channel to see if it is busy. If busy, the client waits a random amount of time before attempting to grab control again. These unpredictable wait periods adversely affect the performance of time-sensitive applications, such as voice communications.

To combat this problem, researchers have proposed several fair scheduling algorithms [15, 16, 22]. The flavor of these algorithms is as follows: the AP is provided the exact state (e.g., number of packets in the queue, their deadlines and priorities) of every client it is servicing; it then schedules the channel for each client in a timely manner. The client state information and the AP schedules are sent over the same data channel.

While being better than IEEE 802.11, these approaches still have some inefficiencies. Since high priority control information has to contend for the channel in the same manner as regular data traffic, undesirable delays are unavoidable.

The multi-radio approach is slightly different. By using the services of the second radio operating on a different frequency (*Design for Choice*), the system parallelizes the scheduling and transmission operations (*Design for Separation*). Not only does this improve timeliness of data transmissions, but also as discussed in Section 4.1.2, it results in energy conservation for clients as they now know exactly when they are scheduled to receive packets or when they are scheduled to get the channel for transmission.

5. CONCLUSION

The users of wireless networks continue to demand better performance and more features. We believe that the current single radio wireless platform is incapable of meeting these growing demands. Despite significant ongoing research, no single technology, including the much-hyped ultra-wide band (UWB) standard [4], has emerged that allows a single radio to deliver high data rates at long range, while consuming low power. Some of the reasons behind this failure are regulatory, while others are more fundamental, grounded in the laws of Physics.

In this paper, we have made a case for re-thinking the core design of the current wireless platform. We have presented a new design that includes multiple radios that work together to accomplish a common task. By using multiple radios, each of which does different things well, and by integrating them at the systems level, such a platform will provide improved performance and greater functionality to the users. We hope that this discussion will encourage further work on these ideas from other researchers.

6. REFERENCES

[1] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks. Technical report, July 2003, (also in IEEE BroadNets 2004).

[2] H. Balakrishnan, S. Seshan, and R. Katz. Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks. *ACM Wireless Networks 1995*, 1(4), December 1995.

[3] N. Bambos and J. M. Rulnick. Mobile Power Management for Maximum Battery Life in a Wireless Communications Network. In *IEEE INFOCOM*, March 1996.

[4] F. C. Commission. FCC 02-48: Revision of Part 15 of the Commission's rules regarding Ultra-Wideband transmission systems. http://www.uwb.org/files/new/FCC_RandO.pdf, April 2002.

[5] A. Communications. Three WLAN Standards in a Single Solution - The AR50001X Combo WLAN Solution. http://www.atheros.com/pt/ar5001X.html, 2004.

[6] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. High-throughput path metric for multi-hop wireless routing. In *ACM MOBICOM*, San Diego, CA, September 2003.

[7] R. P. Draves, J. Padhye, and B. D. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *ACM MOBICOM*, Philadelphia, PA, September 2004.

[8] L. M. Feeney and M. Nilsson. Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking Environment. In *IEEE INFOCOM*, 2001.

[9] J. A. Guitierrez et al. IEEE 802.15.4: A Developing Standard for Low-Power Low-Cost Wireless Personal Area Networks. In *IEEE Network Magazine*, September 2001.

[10] H.-Y. Hsieh and R. Sivakumar. A Transport Layer Approach for Achieving Aggregate Bandwidths on Multi-homed Mobile Hosts. In *ACM MOBICOM*, Atlanta, GA, September 2002.

[11] IEEE. IEEE 802.1x-2001 IEEE Standards for Local and Metropolitan Area Networks: Port-Based Network Access Control, 1999.

[12] IEEE802.11b/D3.0. Wireless LAN Medium Access Control(MAC) and Physical (PHY) Layer Specification: High Speed Physical Layer Extensions in the 2.4 GHz Band, 1999.

[13] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad-hoc wireless networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*. Kluwer Academic Publishers, 1996.

[14] R. Krashinsky and H. Balakrishnan. Minimizing Energy for Wireless Web Access with Bounded Slowdown. In *ACM MOBICOM*, September 2002.

[15] T. Nandagopal, S. Lu, and V. Bharghavan. A Unified Architecture for the Design and Evaluation of Wireless Fair Queueing Algorithms. In *ACM MOBICOM*, July 1999.

[16] E. Ng, I. Stoica, and H. Zhang. Packet Fair Queueing Algorithms for Wireless Networks with Location-Dependent Errors. In *INFOCOM*, March 1998.

[17] L. Nord and J. Haartsen. *The Bluetooth Radio Specification and The Bluetooth Baseband Specification*. Bluetooth, 1999-2000.

[18] H. Packard. HP iPAQ h6315 Pocket PC - Phone Edition. http://www.shopping.hp.com, 2004.

[19] V. N. Padmanabhan and R. Caceres. Fast and Scalable Wireless Handoffs in Support of Mobile Internet Audio. *Mobile Networks and Applications*, 3(4):180–188, December 1998.

[20] R. A. Powers. Batteries for Low Power Electronics. In *Proceedings of the IEEE 83(4)*, April 1995.

[21] J. Rabaey et al. PicoRadios for Wireless Sensor Networks: The Next Challenge in Ultra-Low-Power Design. In *Proceedings of the ISSCC*, February 2002.

[22] P. Ramanathan and P. Agrawal. Adapting Packet Fair Queueing Algorithms to Wireless Networks. In *ACM MOBICOM*, 1998.

[23] RF Monolithics. TR1000/TR1100 916.50 MHz Hybrid Transceivers. http://www.rfm.com.

[24] C. E. Shannon. A Mathematical Theory of Communications. *Bell System Technical Journal*, 27:379–423 and 623–656, July and October 1948.

[25] E. Shih, P. Bahl, and M. Sinclair. Wake On Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices. In *ACM MOBICOM*, September 2002.

[26] S. Singh and C. Raghavendra. PAMAS: Power Aware Multi-Access protocol with Signalling for Ad Hoc Networks. *ACM Computer Communications Review*, 28(3), July 1998.

[27] The Network Simulator: NS-2. http://www.isi.edu/nsnam/ns.