SwitchR: Reducing System Power Consumption in a Multi-Client, Multi-Radio Environment

Yuvraj Agarwal
CSE, UC San Diego
yuvraj@cs.ucsd.edu

Trevor Pering
Intel Research
trevor.pering@intel.com

Roy Want
Intel Research
roy.want@intel.com

Rajesh Gupta
CSE, UC San Diego
gupta@cs.ucsd.edu

Abstract

Multiple wireless network interfaces in a single mobile device exist in order to support their diverse communications and networking needs. This paper proposes a general switching architecture, SwitchR, for managing radio communications for multiple (client) devices utilizing multiple heterogeneous radios per device. SwitchR is deployable incrementally within existing wireless infrastructures, and considers the load imposed on the wireless channel by other communicating clients. SwitchR demonstrates reduction in energy consumption of a mobile device by 47% - 72%, depending upon the application, over the Power Save Mode in WiFi and 13% - 60% reduction in energy over previous multi-radio architectures that do not consider the interactions between multiple clients.

1. Introduction and Motivation

Wearable and mobile devices increasingly feature multiple radio technologies such as cellular, wireless LAN and personal area networks in response to the increasing capabilities of such computing platforms. These *heterogeneous* radios present diverse capabilities, in terms of range of operation, nominal bandwidth, latency and power consumption characteristics. In most platforms, the radio subsystems – whether the RF electronics or transmitted power – constitute up to 50% (786 mW WiFi, 81mW BT out of a total of 1.3 W for the device with the LCD turned off [9]**)** of the total platform power [2][9][10].

Based on their origins, each of these radios have been architected for a specific purpose. As a consequence, these radios and their network interfaces are optimized to provide different forms of energy efficiency, depending on their primary design target. For short distances and low bandwidth connections, Bluetooth is highly efficient consuming on the order of 70 mW for active transfers, compared to almost 800 mW for active WiFi radios [9]. Yet, for high-throughput applications, WiFi provides a lower energy/bit interface at 0.14mW/kbps compared to >0.22 mW/kbps for Bluetooth. Therefore for high throughput applications WiFi is more energy efficient than Bluetooth, which is more suited for lower data-rate or long idle conditions. This paper presents a major generalization of earlier work by exploiting knowledge of network-level parameters and the application needs at individual nodes when making switching decisions. First, it considers the traffic patterns of multiple communicating clients in a network. Second, it enables incremental deployment into already existing wireless infrastructures.

This paper is primarily a demonstration of the fact that the envisioned ‘multi-client switching policy,’ using both local and global channel information, leads to much more energy efficient switching decisions than can be taken by clients independently. This paper makes two primary contributions towards an effective multi-client multi-radio switching system:

- The SwitchR energy-saving switching architecture, which utilizes *independent* low-power Blueooth enabled APs that are incrementally deployable within an *existing* WiFi infrastructure.
- A multi-client switching policy that enables energy efficient communication and networking among multiple simultaneously communicating clients within a multi-radio environment.

2. Related Work

Several techniques have been proposed to wireless power consumption beyond simple *idle* power-save modes. For instance, explorations of systems based on a single available radio address the range from protocol optimizations at the application layer [5][8] to those at the transport layer [4]. Optimizations at the MAC layer [7][11] usually adjust tunable parameters of the 802.11 Power Save Mode (PSM) [6]. However, these optimizations are limited to the energy savings enabled by the low power mode of the WiFi radio, which is still substantial compared to other low power radios such as Bluetooth.

Multiple-radio systems leverage the availability of multiple radios on the same device. Cell2Notify [2] and Wake-on-Wireless [10] investigate the use of a second low power radio purely for wake-up purposes. However, these approaches do not take advantage of

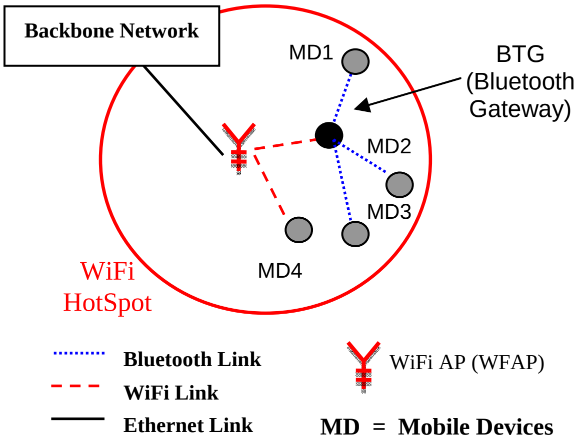

Figure 1: System Architecture – Scenario with multiple MDs, with the BTG node serving as the Bluetooth Gateway. MD1 – MD4 can connect through the BTG using Bluetooth or directly to the existing WiFi-AP using WiFi. For example, MD4 is directly communicating with the WiFi AP.

the fact that the lower power radio can be used for active data transfer as well. *data-on-lpr* [3], proposes the use of a low power radio for active data transfer however, does not evaluate the benefits of such a scheme. Our earlier work on the CoolSpots project [9] uses both Bluetooth and WiFi to optimize the power consumption of a *single* client, using a unified Bluetooth and WiFi access point. The SwitchR system differs from CoolSpots in two main respects: First, the switching architecture does not require unified APs and thus can be easily deployed within *existing* WiFi infrastructure. Second, the SwitchR switching policies consider the interactions between multiple wireless clients.

3. System Overview

The SwitchR architecture, shown in Figure 1, introduces a low-power Bluetooth Gateway (BTG) device into already existing WiFi infrastructure networks. The BTG utilizes the Bluetooth PAN profile to provide network layer (IP) connectivity to other Bluetooth devices. An unmodified WiFi AP is connected to the backbone network over an Ethernet link, while the BTG can be connected to the backbone network either over Ethernet or over WiFi.

While a mobile device (MD) is communicating using Bluetooth, its network traffic is routed through the BTG, allowing the MDs WiFi interface to be switched off. Subsequently, when an application executing on the mobile device requires a higher-bandwidth connection, the MD can turn on its 802.11 interface and access the infrastructure’s WiFi APs directly. Switching between the WFAP and BTG in the SwitchR architecture is accomplished by network level reconfiguration using Address Resolution Protocol (ARP) adjustments in the network and route-table updates on the MD as well as the BTG. Further details about the switching mechanism used in SwitchR are available in [1].

The decision on *when* to switch from one radio to the other is governed by switching policies. In the case of multiple communicating MDs, the policies for switching between various interfaces must take into account the dynamic nature of the Bluetooth channel as the presence of other MDs affects the total bandwidth available, in addition to the link quality of the Bluetooth channel. A hybrid approach that takes into account both the MDs application requirements and the effective capacity of the wireless channel is thus needed to design effective switching policies.

4. Experimental Setup

An experimental test bed (Figure 1) consisting of multiple wireless nodes placed at various fixed locations in a moderately sized laboratory (8m by 12m) is used to test the SwitchR framework [1]. We use a Linksys router (BEFW11S4) as our WFAP and have a dedicated Test Machine (TM) to generate test traffic. The MDs and the BTG are based on the Stargate2 research platform, which has an on-board Bluetooth radio (Bluecore3) and a compact flash based WiFi radio (Netgear MA701). Each MD is instrumented with an integrated power measurement capability and also monitors its own network traffic to log the amount of data transferred. Using this distributed power measurement and data logging capability, we can simultaneously measure the energy consumption for all of the test devices to get a detailed characterization of the overall system power consumption.

Our experimental design consists of four benchmark tests running on the four MDs; where in any run each executes a different benchmark. We ensure that each benchmark executes at least once on each device, factoring out any hardware variance between individual devices. In any run, all devices use the same policy; each benchmark suite is replicated for each of the four policies, resulting in 4 (benchmarks) x 4 (devices) x 4 (policies) = 64 benchmark runs for a set of results. The benchmark themselves execute in a continuous loop (since they are not necessarily the same length), and an individual result consists of a fixed-length sample of different statistics (e.g. power consumed) consisting of at least two complete benchmark executions.

5. Benchmarks

Since our evaluation focuses on a multi-client scenario, we use a set of *n* benchmarks to constitute an *application suite*, where *n* corresponds to the number of MDs in our test setup. The baseline benchmarks we use are the *idle* and the *transfer* benchmarks. The *idle* benchmark is the state of the system when there is no data transfer taking place, while the *transfer* benchmark represents a TCP stream that tries to send data as fast as it can.

The *streaming* benchmark models viewing live video content or streaming audio MP3s. These applications have real-time requirements and need QoS guarantees. Most media streams are sent over UDP and the two QoS metrics that are often used are *jitter* and *packet loss*. The standard *Iperf* tool is used to generate various sets of traffic patterns. We use three streaming benchmarks: stream128, stream156, and g711 (VoIP codec) with data rates of 128Kbps, 156Kbps and 64Kbps respectively. These sample bit rates can be handled by our current BT v1.2 hardware (1Mbps).

The *web* benchmark emulates the traffic pattern of a web browsing session. We monitored the web browsing traffic of a typical user and then downloaded the content that they visited locally. In addition, we measure the inter-arrival time between subsequent page requests capturing the user "think" time. Our goal in creating this benchmark was to emulate a session with sporadic data transfer characteristics, (small periods of activity interspersed with idle intervals), and measure its corresponding effect on the other benchmarks.

6. Switching Policies

The two main switching policy decisions to be made are: (a) When to switch on the high power, high throughout (WiFi) radio, and (b) when to switch back down to the low power, low throughput (Bluetooth) radio. Excessive switching can potentially increase power consumption and adversely affect applications. Optimized switching is necessary to realize the potential in energy savings afforded by multiple radios.

6.1 Baseline Policies

The *wifi-CAM, wifi-PSM* policies serve as baselines for evaluating the energy and performance behavior of the system. *wifi-CAM,* operates the WiFi radio in always-on mode. *wifi-PSM* and all the other policies use the Power Save Mode (PSM) of WiFi [6], which essentially duty cycles the WiFi radio.

6.2 Cap-Dynamic Policy

The *cap-dynamic* policy was the most energy-efficient policy from CoolSpots [9], which looked at the current capacity of the Bluetooth channel to make switching decisions. It works well for single client situations, however in multi-client situations it has significant problems correctly predicting the available capacity since it does not account for other simultaneously communicating clients. Details and discussion about the *cap-dynamic* policy are shown in [1][9].

6.3 Multi-Client Policy

The *multi-client* policy takes a different approach to determine the appropriate switching points. For switching-up to WiFi it uses a combination of multiple echo-response packets and the Received Signal Strength Indication (RSSI) of the BT link to estimate channel quality. If the average RSSI of the BT link degrades, and/or the echo-responses time increases substantially it signals a drop in channel quality and a switch-up to WiFi is triggered.

The switch-down case to Bluetooth is a combined decision that involves the MD as well as the Bluetooth Gateway. At the BTG the maximum bandwidth $MAXBW_{bt}$ that the BT interface can support is first estimated empirically (450Kbps for our setup). For switching-down the policy (executing on the MD) periodically measures the average bandwidth on the WiFi channel. If the average bandwidth observed on the WiFi interface is greater than $MAXBW_{bt}$ then the policy reverts back to measuring the WiFi channel as there is no point in switching down to BT given the current application requirements. However if the bandwidth measured on WiFi is less than $MAXBW_{bt}$, then the *multi-client* policy performs multiple checks to determine whether it is optimal to switch down to Bluetooth.

First the policy estimates the quality of the Bluetooth link by measuring the RSSI and the time for multiple echo-response packets. If channel quality is good, the *multi-client* policy queries the BTG and sends the average application bandwidth requirements as measured on the WiFi channel. The BTG continuously measures the bandwidth it observes through its BT interface and in case there is some spare capacity (bandwidth < $MAXBW_{bt}$) it sends a confirmation back to the particular MD that sent the query. If the MD receives a positive conformation it switches down to BT, otherwise reverting back to this decision process.

7. Results

Figure 2 summarizes the impact of each policy for two separate benchmark suites. Figure 2a considers the four basic benchmark types, and highlights the overall effectiveness of the multi-radio switching concept. Figure 2b considers a more loaded scenario that stresses the capacity of the underlying Bluetooth channel, highlighting the changes introduced in the multi-client policy. The overall results are not surprising: Idle shows great savings, transfer shows very little savings, and the streaming media and web benchmarks show varied savings depending on context. As illustrated in Figure 2b, the multi-client policy saves up to

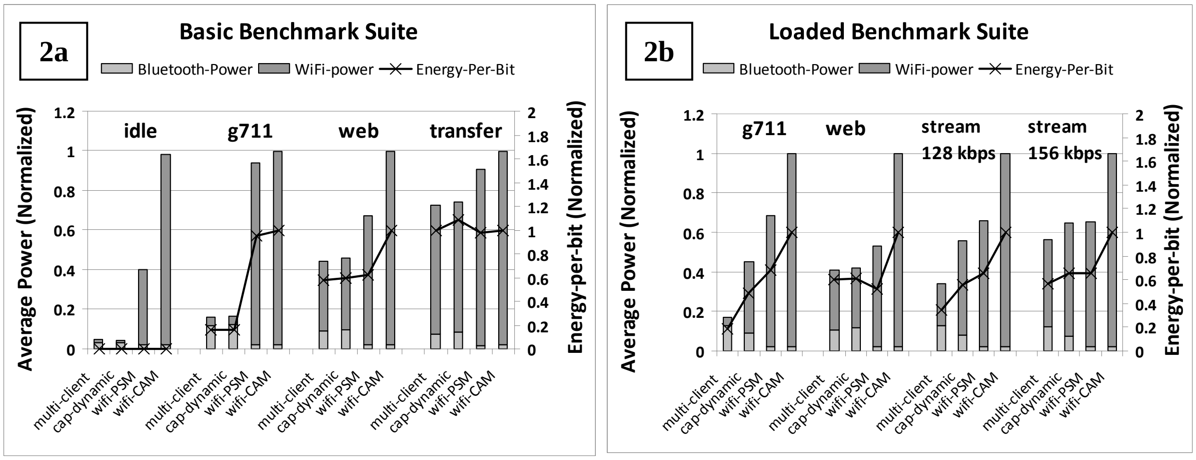

Figure 2: Results for the switching policies for two representative benchmark suites. These graphs show how, in some circumstances, the multi-client policy dynamically adapts to the changing conditions of the wireless channel. (Note: Each bar represents an average of 4 runs for each benchmark)

62% over the cap-dynamic policy and up to 72% energy over the wifi-PSM, depending on the application.

The multi-client policy shows its main improvement for the streaming media benchmarks, as shown in Figure 2b. These workloads are relatively constant, and the corresponding switching decision is dictated primarily by the behavior of the *other* nodes in the system (e.g., a change in workload by the web benchmark). The primary drawback with the cap-dynamic policy is that it only considers the data traffic through the respective device itself, and ignores other traffic on the wireless channel: when the web benchmark stops transferring data, the cap-dynamic policy does not adjust to make use of the now free Bluetooth channel.

8. Conclusions

In this paper we have presented *SwitchR*, a novel multiple-radio based switching architecture, enabling wireless devices to use standard wireless applications yet significantly increase their battery operating time. A major advantage of our SwitchR architecture is that it is incrementally deployable within existing WiFi infrastructure, and that it can be used without modifying client applications. Furthermore, SwitchR performs well even with *multiple* simultaneous communicating clients, and reduces the energy requirements of all participating devices substantially. For our suite of representative benchmark applications, the multi-client policy enables energy savings up to 72% over the WiFi Power Save Mode (PSM), and up to 60% compared to previous multi-radio architectures. Such low power operation is attractive to both traditional mobile technologies, as well as emerging wearable systems, which are both highly power-sensitive devices.

9. References

[1] Agarwal, Y., Pering T., Want R., Gupta R. "SwitchR-Reducing System Power Consumption in a Multi-Client, Multi-Radio Environment", *UCSD CSE Technical Report*, CS2008-0926

[2] Agarwal, Y., Chandra R., Wolman A., Bahl P., Chin K., Gupta R., "Wireless Wakeups Revisited : Energy Management for VOIP over Wi-Fi Smartphones", In *Proc. ACM MobiSys* ,2007

[3] Bahl, P., Adya, A., Padhye, J., and Wolman A., "Reconsidering Wireless Systems with Multiple Radios". *In ACM SIGCOMM Computer Communications Review (CCR)*, 2004

[4] Bertozzi, D., Raghunathan, A., Benini, L., and Ravi, S.. "Transport Protocol Optimizations for Energy Efficient Wireless Systems". In *Proc. of DATE*, 2003.

[5] Flinn, J., and Satyanarayanan M., "Managing Battery Lifetime with Energy-aware Adaptation". In *ACM Transactions on Computer Systems (TOCS)*, May 2004.

[6] IEEE 802.1x-2001. IEEE Standards for Local and Metropolitan Area Networks, 1999

[7] Krashinsky, R., and Balakrishnan H., "Minimizing Energy for Wireless Web Access with Bounded Slowdown", In *Proc. ACM Mobicom*, 2002

[8] Kravets, R. and Krishnan P., "Application-driven Power Management for Mobile Communication," *Wireless Networks*, vol. 6, no. 4, 2000

[9] Pering, T., Agarwal, Y., Gupta R., Want R., "*Coolspots:* Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces". In *Proc. ACM MobiSys*, 2006

[10] Shih, E., Bahl P., and Sinclair, M.J., "Wake on Wireless: an Event Driven Energy Saving Strategy for Battery Operated Devices". In *Proc. ACM Mobicom*, 2002.

[11] Woesner, H., Ebert J.P., Schlager M., and Wolisz A., "Power saving mechanisms in emerging standards for wireless LANs: The MAC level perspective". In *IEEE Personal Communications*, vol. 5, pp. 40–48, June 1998.