Ph.D. Thesis Proposal

Open Networking Infrastructure: Boosting Wireless Networks in the Era of Cloud

Tao Jin

COLLEGE OF COMPUTER AND INFORMATION SCIENCE

NORTHEASTERN UNIVERSITY

Ph.D. Committee

Advisor Guevara Noubir

Alan Mislove

Rajmohan Rajaraman

Ext. member Prasant Mohapatra **University of California, Davis**

August 31, 2012

Abstract

In recent years, the mobile Internet underwent revolutionary changes, and has fundamentally changed the way users access the Internet. This resulted in an unprecedented demand for ubiquitous network access. However, due to the fundamental radio frequency communication constraints, the service quality and scalability of cellular systems is limited. Wi-Fi, now being a well-developed standard technology, has been densely deployed, especially in urban area.

My research work is primarily motivated by these exciting trends in mobile networks. We believe WiFi has the unique potential to form a community infrastructure that provides a truly scalable, efficient and ubiquitous access to wireless and data, especially in urban area.

In this work, we propose an *Open Infrastructure* framework. Based upon this framework, we explore a group of mechanisms to boost the mobile computing experience with the leverage of urban WiFi. To evaluate the feasibility of our ideas, we have built and deployed an Open Infrastructure testbed, consisting of 30 customized home WiFi APs running in the urban areas of Boston and Houston. Since February 2011, we have collected over 70 million residential network usage statistics record and 1.3TB of traffic data trace. This testbed provides us with first-hand information on urban WiFi and a realistic setup to try out a variety of research ideas.

Our research has been focusing on two domains, urban WiFi assisted energy saving on mobile devices and idle bandwidth harvesting from home WiFi APs. We have initiated two projects, WiZi-Cloud and BaPu, to develop a set of enabling mechanism and prototypes to demonstrate the feasibility of our ideas.

In WiZi-Cloud, we extend the current WiFi AP and mobile device with an alternative ultra-low power ZigBee radio interface, to reduce the battery consumption on the energy constrained mobile devices. The WiZi architecture and its support for multiple heterogeneous radios is transparent to the applications and allows seamless interfaces switching. WiZi supports multiple ZigBee transceivers and channel coding schemes. We design and prototype a complete suite of hardware/software solution. Our experimental results show that the WiZi-Cloud well supports a large set of mainstream mobile applications and improves the energy efficiency by 3 folds and can exceed WiFi coverage range.

BaPu is motivated by the fact that today's residential broadband connections have limited backhaul bandwidth, especially in uplink, which highly constrains many fastly growing applications, such as HD content instant sharing, efficient cloud storage backup, etc. In BaPu, we design the mechanisms of aggregating the idle broadband uplinks, by having the mobile device communicate with multiple proximate WiFi APs in the same neighbourhood. BaPu is a complete software solution running on home WiFi APs, and requires no modifications to the client devices. This makes an easy incremental adoption of BaPu technology. Our architecture and protocols design provide a clean solution to several challenges in efficiently supporting both TCP and UDP. Our prototype system shows that BaPu can efficiently aggregate the backhaul bandwidth of multiple access points and achieves up to 90% of the theoretical maximum throughput.

Contents

1 Introduction

1.1 Fast Evolving Mobile Networks

In recent years, the mobile Internet underwent revolutionary changes. Mobile devices show a very robust and rapid growth. According to a press release from the research firm Canalys [1], in 2011 the annual global shipments of smart phones exceeded those of client PCs for the first time. From 2011 to 2016, smart phone unit sales are predicted to grow at about 28 percent compound annual growth rate [2]. The growth of the mobile Internet is stimulated by the fast expansion of mobile data networks, such as 3G and 4G networks, as well as the increasingly powerful mobile devices hardware.

This mobile revolution has fundamentally changed the way users access the Internet. The growing bandwidth of mobile networks, combined with the increasingly powerful mobile devices, are making a large variety of mobile applications more feasible. Such trend also nurtures end users' ever increasing demand of accessing the content of their interest on the go. Mobiles devices, such as smart phones and Internet connected tablets, are currently the key platform for our work, entertainment, social activity, and have become an indispensable part of our daily life. End users download a large amount of content from the Internet Cloud to their devices, to browse news, to stream music and videos, and to keep in touch with their friends and colleagues. However, in the era of Cloud, the mobile application is far beyond only fetching content from the Internet. In recent years, the User Generated Content (UGC) is also experiencing an explosive growth. The mobile devices are equipped with high-resolution cameras and a variety of sensors, and are quickly becoming the primary devices to generate personal multimedia content. End users have significant demand of sharing these high volume of UGC with others in an instantaneous way, or backing them up in the "Cloud Storage". According to recent reports, 250 million photos are uploaded to Facebook every day, over 72 hours videos are uploaded to YouTube every minute, and 325 million files are saved on Dropbox every day [3–5]. IDC [6] forecasts that by 2015, more U.S. Internet users will access the Internet through mobile devices than through PCs or other wireline devices. Global mobile data traffic is set to increase 18-fold between 2011 and 2016 [7].

1.2 New Challenges

Such trend of ubiquitous communication allows users to do much more on the go, which drives an unprecedented demand for bandwidth. Therefore, cellular networks are stretched to their limits when serving rich content to a large number of users. In some urban areas, dropped calls can reach 30% [8–10]. The service quality and scalability of cellular systems is limited by fundamental constraints. Even if 3G and 4G systems, such as Long Term Evolution (LTE) and WiMax, can provide data rates of tens of megabits per second (e.g., the LTE specification provides 100 Mbps peak downlink and 50 Mbps peak uplink for a 5MHz RF band), this is shared among all the users of a base station. Therefore, scaling cellular networks requires a high density of base stations [11]. This results in a substantial cost in terms of sites construction, backbone connectivity, and maintenance (or necessitates the cooperation of users to deploy femto base-stations in their homes), in addition to the cost of leasing the cellular RF spectrum. AT&T, for example, needs to spend tens of billions of dollars every year to upgrade their wireless networks to sustain the fast growing mobile traffic [12]. Meanwhile, the flat-rate, unlimited cellular data plans continuously deflate the mobile operator margins and place a lot of stress on mobile operators' revenue model, with revenues dropping from US$0.43 per MB to approximately $0.02 per MB in 2014, according to a 2010 study by Bernstern Research [13].

Therefore, driven by the users' demand, business model, technology limitations in current mobile networks, and especially the traffic surge in mobile networks [14], the mobile operators have been showing significant interest in adopting heterogeneous networking technologies to offload the traffic pressure from cellular data networks to augment both coverage and capacity. Two primary offload technologies used by the industry are Femtocells and WiFi [15–19].

1.3 Data Offload Technologies

Femtocell is a small cellular base station, designed primarily for user in home and small business. Femtocell uses standard cellular radio technology, and is typically connected to the wired backhaul, such as cable or DSL. Femtocells allows ISPs to extend the cellular coverage indoor or at some place with limited or unavailable cellular service. The major advantage of femtocell is that any cellular mobile devices can function with femtocells, and can seamlessly roam between femtocells and regular base stations. However, since femtocells operate within licensed spectrum which needs careful planning, the unplanned femtocells deployment may lead to severe interference issue and thus inefficient spectrum utilization.

WiFi now being a well-developed standard technology, has been densely deployed, especially in urban area. Since WiFi is designed to operate within unlicensed spectrum (i.e., 80 MHz in the 2.4GHz ISM band, and 240 MHz in the 5.15-5.824 GHz U-NII bands), how to alleviate the inter-channel interference has already been taken into account. Besides, since the unlicensed spectrum is far greater than the licensed cellular spectrum, more traffic can be delivered using WiFi. Thus, WiFi has the characteristics to form the building blocks of seamless, ubiquitous, and eternal wireless access. Over the years, WiFi are very fragmented. Recently, there are two main approaches to consolidate WiFi networks, community approach and carrier approach. One typical example for the former approach is FON [20]. By provideing customized consumer WiFi access points and software solutions running on FON's backend servers, FON allows end users to share part of their bandwidth as public WiFi signal for free or for some payment in return. FON claims to have over 5 million WiFi hotspots over the world. In the latter case, mobile operators deploy and manage a large number of carrier level WiFi hotspots to offload the data from the congested cellular infrastructure. By recent research reports from Infonetics [21, 22], 79% ISPs will deploy WiFi hotspots by 2013 and a significant growth of carrier deployed WiFi is expected in the coming year. Asides from augmenting mobile services, ISPs are also seeking to make WiFi become an integrated part of mobile network. An amendment to the IEEE 802.11 standard, 802.11u, is being developed to make the WiFi data offload more efficient and seamless [23–25].

My research work is primarily motivated by such exciting trend in mobile networks. Given the characteristics of WiFi presented in the above section, we believe WiFi has the unique potential to form a community infrastructure that provides a truly scalable, efficient and ubiquitous access to wireless and data, especially in urban area. Our project is founded upon our recent work on analysing and characterizing residential broadband traffic, inter-AP connectivity, intra-ISP vs. inter-ISP APs connectivity, and AP-to-CDN connectivity. With our Open Infrastructure testbed [26], which consists of 30 highly customized home WiFi APs deployed in urban area in Boston and Houston, we have been monitoring the residential broadband usage for a 18-month period since February 2011. During this period, we have collected over 70 million residential network usage statistics record and 1.3TB of traffic data trace. Besides, to have a better understanding of residential WiFi from a large scale, we carried out Wardriving [27] experiment in Greater Boston Area and characterized tens of thousands of urban WiFi APs and their Internet connectivity. Our results show that the residential broadband resource is under-utilized, and the high-density ready-to-use urban WiFi generally has good connectivity to Internet. Based upon our preliminary measurement results, we proposed a research framework called Open Infrastructure. Such framework opens up a broad range of research topics. Our research goal is to develop enabling mechanisms, protocols and demonstrators to leverage the urban WiFi to provide a truly scalable, secure, efficient and ubiquitous access to wireless and data.

In section 2, we will present in detail our proposed Open Infrastructure framework, and discuss the potential research topics based upon such framework. In section 3, we will describe the Open Infrastructure testbed we have built and deployed in the urban area in Boston and Houston, and our preliminary measurement results. Among a variety of potential research topics of Open Infrastructure, our research has been focusing on two aspects, energy efficiency and idle bandwidth aggregation. We have initiated two projects, WiZi-Cloud and BaPu, on these two topics. In section 4 and 5, we will present these two projects in detail. In section 6, we will present the future task items, and my tentative task schedule.

2 Proposed Research

As we discussed in previous section, the mobile operators have been showing significant interest in adopting WiFi technology as part of mobile strategy. In this section, we first briefly describe the current attempts made by the mobile operators to leverage the urban WiFi infrastructure to make it an integral part of the carrier's networks. However, we believe the urban WiFi will have a far greater potential than just providing a point of access to the Internet. We will present our vision and our proposed *Open Infrastructure* framework.

2.1 Industry's Point of View

Along with the great success of smartphones such as iPhone and Android, the mobile operators' business models are transitioning from mobile voice model to mobile data model. The operators start paying more attention to how to make better use of the unlicensed spectrum and WiFi as their mobile strategy. WiFi, now being wildly adopted worldwide, is very fragmented. WiFi are mostly deployed for use in home and enterprise environment, with loose management and limited or no interactivity. Although the mobile operators are installing carrier level hotspots, each operator manages their hotspots separately, and serve isolated user base. This results in a lot of usability issues in today's hotspots. For example, users must maintain separate credentials to authenticate with different hotspots; roaming between hotspots managed by different operators is limited; public hotspots are not encrypted, and thus mobile users are exposed to serious security threats; when multiple SSIDs are present in the same location, which (optimal) SSID should the mobile device be associated to is not an automated process. Lacking sufficient usability, today's hotspots may place a lot of constraints on the services with mobility requirement. Therefore, the mobile operators and WiFi product vendors are currently focusing on standardizing a new set of technologies, such as 802.11u [28], Hotspot 2.0 [23, 29], etc., to improve the mobility, security and usability problems in today's WiFi infrastructure. Besides, they are also actively promoting new technologies, such as 802.11ac [30] to make the ultra high bandwidth available in their WiFi networks.

2.2 Our Vision - Open Infrastructure Framework

The active efforts in industry help improve the accessibility. However, as we take a deeper look at this problem, we believe that the urban WiFi infrastructure can play an important role, that is far beyond a simple access point. Each WiFi AP is connected to the Internet through a broadband connection, which is often times under-utilized. Besides, as the hardware becomes commodity, today's WiFi APs are equipped with faster CPU and more RAM, which allows the APs to do much more than they do today. We proposed an framework, called Open Infrastructure.

Open Infrastructure is an urban area WiFi infrastructure instrumented with customized computation, communication and storage capability to provide users with a scalable, secure, efficient and ubiquitous access to wireless and data. Figure 1 depicts the Open Infrastructure architecture, and illustrates a set of new potentials of ubiquitous computing with the leverage of urban WiFi. On one hand, by turning WiFi from a loosely managed fragmented infrastructure to a well managed, tightly coordinate one, we expect to provide much better Internet accessibility from a variety of perspectives. Current hotspots are all unsecured due to the difficulty of key distribution in advance, which poses security thread while users access networks through public hotspots. Some new WiFi authentication mechanisms (*Case 1*), such as [28, 31], can make the WiFi hotspot access a more transparent and seamless process. Besides, in the community based WiFi infrastructure, how to provide appropriate incentives for both regular users and carriers to participate in scaling community networks and contributing idle resources can lead to a group of game-theoretic and mechanism design formulations [32–38]. As the WiFi infrastructure keeps scaling, the cross-channel interference issue becomes more and more serious and limits the performance. Some much finer spectrum management (*Case 4*) than existing CSMA/CA is necessary for a truly scalable solution. Currently there is a trend of offloading the spectrum management and coordination to the cloud [39–45]. In addition, the growth of WiFi bandwidth is outpacing the growth of residential broadband. The latest 802.11ac [30] can provide up to 7Gbps, which is far greater than today's typical cable connection (20Mbps for downlink, 2-3Mbps for uplink). This leads to a group of study of examing the mechanisms to aggregate the idle bandwidth behind the WiFi APs (*Case 6*) to boost the downlink and uplink performance [46–54]. With fast enough Internet connectivity through WiFi APs, in addition to traffic loading from cellular networks to WiFi, a lot of computation work can also be offloaded to the remote server farms in the cloud in order to reduce the battery usage on the energy constraint mobile devices (*Case 2*). Recent studies have shown the feasibility and benefit of such mechanisms [55]. Furthermore, as the hardware becomes commodity, the WiFi APs can easily be extended with storage capability. Such large set of storage units, residing at the edge of the cloud, can form a distributed storage service (*Case 3*). Some of recent studies [56–59] examine the feasibility and advantages of such system compared with current main stream consolidated storage services. Asides from providing the online storage service to end users, the WiFi APs can also function as a Content Delivery Network (CDN) (*Case 5*), because urban WiFi APs are generally closer to the user and consume little energy, and thus have the potential of forming a P2P style CDN for better throughput to users, less operational cost, and less energy cost [34, 38, 56–66].

Even though a lot of efforts are being made by industry and academia to leverage the urban WiFi, the mechanisms discussed above may not all be applicable in real world without appropriate business models. Our research is focusing on evaluating the feasibility of such ideas. Based on such architectural design, we have built and deployed an Open Infrastructure testbed comprising of customized home WiFi access points. Based upon this testbed, our research has been focusing on two domains, energy efficiency and idle bandwidth harvesting. We have initiated two projects, WiZi- Cloud and BaPu, to develop a set of enabling mechanism and prototypes to show the feasibility of our ideas:

- **Open Infrastructure Testbed:** currently consists of 30 home WiFi APs customized with a set of software/hardware developed in our lab. The purpose of such testbed is to demonstrate the feasibility of the proposed mechanisms, identify practical implications of the architectural design choices, and enable extensive evaluation of the proposed research framework in an environment that is closest to reality. Such testbed forms an urban WiFi infrastructure and provides the researchers with the first hand information of the urban WiFi networks and a realistic setup to try out a variety of research ideas.
- **WiZi-Cloud**: To explore the possibility of improving the energy efficiency of mobile devices with the assistance of WiFi infrastructure, we proposed WiZi-Cloud solution. Nowadays the mobile devices are quickly becoming the major way of communicating with others. However, to keep the users connected to the Internet, the network interfaces, such as 3G and Wi-Fi, are generally turned on. Our study and the literature works have shown that 3G and Wi-Fi interfaces have become one of the major energy consuming components on the devices, which result in very short battery lifetime and negative user experience. In WiZi-Cloud, we integrate an ultra low power alternative radio to the mobile phone

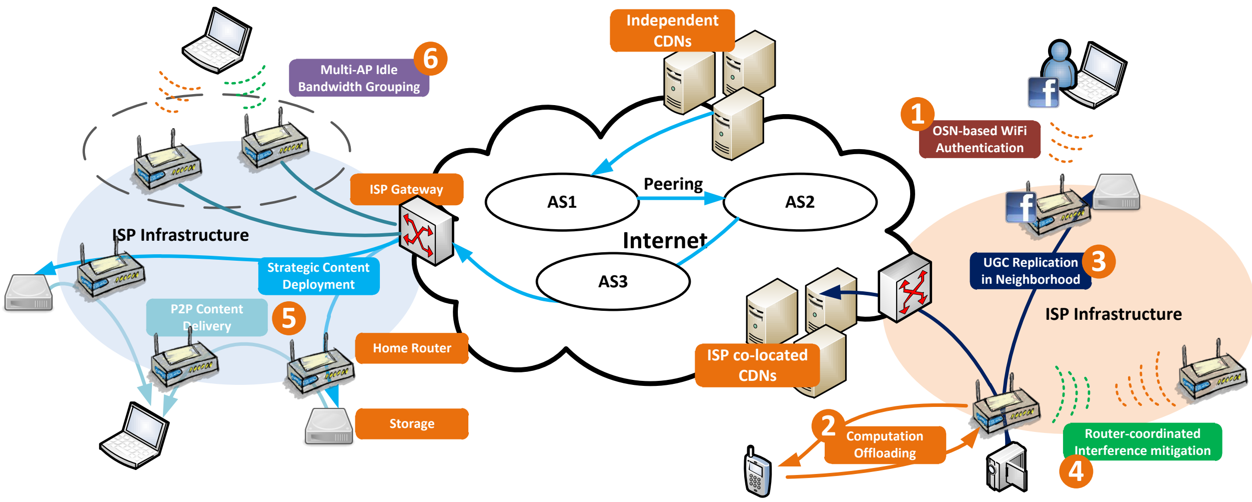

Figure 1: Open Infrastructure Architecture.

and the wireless access point, so that we can use the lower power radio to carry some network traffic when possible, and keep 3G and Wi-Fi off at most time. We implemented a prototype WiZi-Cloud system. Our experiment results showed that the mobile phone standby life time can be extended up to 3 times, and 2 times when the phone is actively transmitting data.

- **BaPu**: Nowadays, the ISP generally offers 10 to 20 Mbps downlink bandwidth, which is sufficient for most Internet applications. In the contrast, the uplink bandwidth is much lower, generally around 1 to 3 Mbps. However, with more powerful handheld devices become available, the amount of User Generated Content (UGC) is rapidly increasing. Users generate more and more HD video clips and pictures, and would like to instantly share them with their friends through Facebook, Youtube, etc. Due to the limited uplink bandwidth, the sharing process is often times unpleasant and infeasible. We are aimed to provide a solution to tackle this asymmetric bandwidth issue in the residential networks and boost the uplink to allow efficient and instant HD content sharing. We designed and implemented BaPu, which is a system running on Wi-Fi APs and application servers. BaPu allows the client device to upload data through multiple APs so that the uplink process aggregates the uplink bandwidth from all APs. Our prototype system shows that BaPu can efficiently aggregate the uplink bandwidth and achieve up to 90% of the theoretical maximum throughput.

3 Open Infrastructure

3.1 Testbed Description

Open Infrastructure testbed is comprised of a set of off-the-shelf home WiFi access points running customized firmware, a back-end server, optional hardware peripherals, and a suite of control and management tools. Our AP firmware is developed based on the popular open source embedded Linux distribution, OpenWRT [67]. OpenWRT supports a large number of APs or embedded products [68]. Our default deloyed APs are Buffalo WZR-HP-G300NH [69], which has a 400MHz CPU, 64MB RAM, 32MB flash, 5 Gigabit ethernet ports and one USB interface. Its WiFi interface consists of the AR9132/AR9103 Atheros Network Processor and Radio chipsets which supports flexible channels, and interference measurements. In our testbed, each AP is extended with either a 16GB USB Flash or a 250GB hard drive (See Fig. 2). The extended storage capability allows the APs to store our experiment results or form a distributed storage system. Since Feb. 2011, we have deployed 30 customized APs in urban areas in Boston and Houston, serving around 100 individual users. Our recruited users are mainly graduate students and professionals of age between 20 and 40, with diverse background. Currently we are expanding the scale of our testbed, aiming to reach 150 nodes in 1 year.

In the customized AP firmware, we have pre-installed a heartbeat client program, which monitors the real-time status of each home network, and periodically reports the network statistics to our back-end server every 10 seconds. The heartbeat report includes the statistics data such as number of connected devices, average bandwidth usage for different traffic types in the past 10 seconds, etc. All the reports are stored in our back-end MySQL server for next step data processing. Since Feb. 2011, we have collected over 70 million records which allows us to monitor the urban network usage in a very long term.

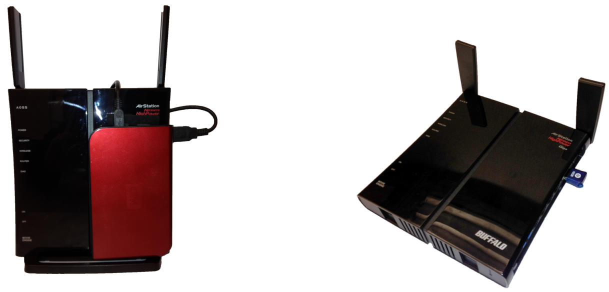

Figure 2: Open Infrastructure Testbed Prototype WiFi AP, Buffalo WZR-HP-G300NH, with 3.5” Hard Drive and 16GB USB Flash

Besides, we developed a suite of SSH-based remote management tools. With such tools, researchers can remotely upgrade the firmware, update the AP configurations, schedule experiment tasks, etc. In order not to disturb our test users’ normal network usage, all APs are scheduled to upload the experiment results to our server during off-peak hours, say midnight during weekdays. Also, we have developed a Web based Testbed Management Portal to better support the management process (See Fig. 3).

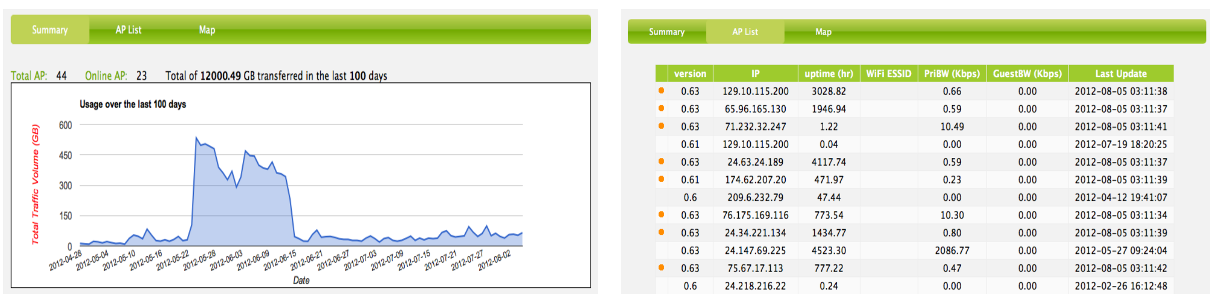

Figure 3: Web-Based Open Infrastructure AP Management Portal

Our major goals of building this testbed are:

- provides the researchers with the first hand information of the urban WiFi network, and the capability of long-term observation.
- Since the urban WiFi networks have unique characteristics compared with academic network and enterprise network, in terms of traffic load, WiFi interference leve, etc. The testbed residing in real neighbourhood provides the researchers with a realistic setup to try out a variety of research ideas targeted at urban network.

3.2 Preliminary Measurement Results

Since our research focus is how to leverage the urban WiFi networks in a community based approach, the actual potential of such crowd sourcing approach largely depends on the density of the urban WiFi, and the amount of resources that each WiFi AP may contribute. Therefore, we would like to first answer two questions: 1) What is the density of current urban WiFi? 2) How much resources we should expect from current urban WiFi? To have a preliminary idea, we carried out a set of experiments on a small deployment of 30 nodes, complemented with extensive wardriving experiments in four residential neighbourhood in Boston area. Our measurement results have several interesting implications as listed below, and provide strong evidence in support of our proposed research:

- **Urban WiFi is generally equipped with under-utilized backhaul broadband**: Such observation is made based on over 700 million per-10 second bandwidth usage records we collected from the Open Infrastructure testbed between February 2011 and May 2012. We count the bandwidth samples in every 1 minute time slots. We count the number of samples that fall into certain bandwidth range as listed in Figure 4, and calculate the probability of bandwidth usage over a certain value at a certain time slot. One immediate observation is that the traffic load in residential networks is rather limited even at the peak hours (between 8pm and 11pm). This applies to both uplink and downlink. For example, during peak hours there is over 65% chance that the residential routers have less that 10Kbps traffic load on both downlink and uplink. This implies we may make better use of the under-utilized backhaul broadband bandwidth.

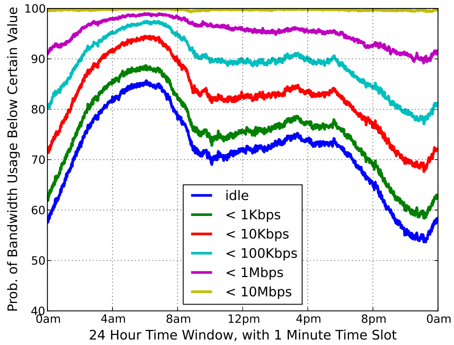

(a) Uplink Bandwidth Usage During 24 Hours

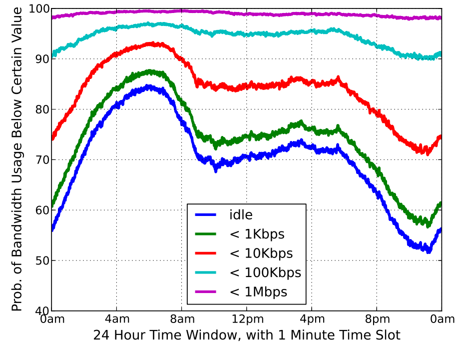

(b) Downlink Bandwidth Usage During 24 Hours

Figure 4: Urban broadband bandwidth usage per minute averaged over 15 months period indicates that the broadband capacity is mostly idle for both uplink and downlink.

- **AP Density in Urban Area:** We carry out a Wardriving experiment in 4 residential areas in Greater Boston, and have collected over 26K Urban WiFi AP information, including ESSID, BSSID, GPS coordinate, signal level, encryption method, etc. There were an average of 17 access points visible in each scanning operation. During the Wardriving, we also attempt to associate to the detected APs with best signal quality. Among all the successful associations, the mean signal quality of the AP reported by the athe9K driver is -81dBm. Take this value as an indicator of a good connection, there were an average of 8 WiFi APs with signal level greater than -81dBm in each scanning operation.
- **RTT implies inter-connectivity:** To examine the latency characteristics in the urban network, we have carried out extensive RTT measurements from our Open Infrastructure testbed to about 1000 major CDN or public web servers. We have also carried out latency measurement during wardriving through public accessible APs. The CDN and web server IP addresses were obtained from our real network traffic data trace collected from our testbed. As shown in Figure 5, within the same ISP, the AP-to-AP round trip time is significantly lower than AP-to-CDN. While 50% of intra-ISP AP-to-AP RTT is below 24ms, 50% of AP-to-CDN RTT are above 46ms. Note that the measurements were done in Boston and Cambridge, MA, the home of Akamai Technologies, one of the largest CDN companies.

  The under-utilized bandwidth, high AP density and short latency between urban WiFi APs confirms the potential for community networks and the necessity of efficient coordination and interference mitigation mechanisms.

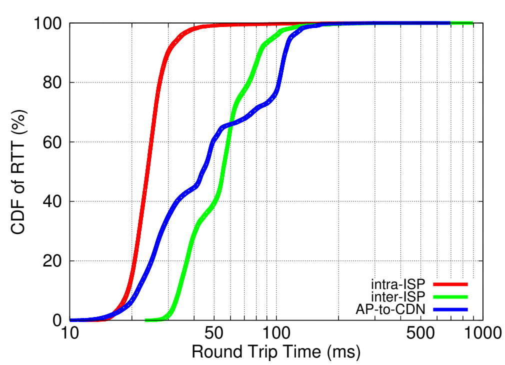

Figure 5: urban WiFi Latency Measurement. The CDF of RTT indicated that the intra-ISP latency is significantly lower than the AP-to-CDN latency.

4 WiZi-Cloud - Improving Energy Consumption by Leveraging Urban WiFi Infrastructure

4.1 Energy Efficiency Issue on Mobile Devices

4.1.1 Mobile Technology Outpacing Battery Technology

Mobile devices have been rapidly evolving in the past few years and their ability has gone far beyond providing telephony services. Nowadays, smartphones are enabling an increasingly large set of applications. A lot of Internet based applications appear on mobile platforms, such as email, VoIP, web browsing, instant messengers and video/audio streaming, and have become more and more popular for daily use. Such applications necessitate a reliable and ubiquitous Internet access, carrying either intensive network traffic or periodic network access. Thus the mobile devices are much reliant on the battery power (See section 4.1.2). Unfortunately, battery technology and mobile technology are at quite different stages. Unlike mobile technology, which is still at a young stage and shows robust and rapid growth, battery technology has been evolving for a century. Battery technology follows the classic "S-curve" of innovation, and shows slow or no improvement since 2005 [70,71]. The energy consumption issue has been one of the major constraints on mobile communications. According to a recent survey, the short battery lifetime is among the top complaints of smartphone users [72].

4.1.2 Energy Consuming Network Interfaces

Smartphones typically access the Internet either through cellular networks or WiFi networks. Due to the fundamental technology constraints, scaling cellular networks requires a high density of base stations which incurs a substantial cost in terms of sites construction and maintenance. WiFi networks can significantly help scale wireless access, in cooperation with cellular technologies, especially within urban areas. WiFi networks have the advantage of operating over large license-free bands, and have been densely deployed in urban areas [73]. In addition, WiFi hardware and standards have been well developed for years. However, it is well known that the WiFi interface on mobile devices suffers from high energy consumption even in Power Saving Mode [74–80].

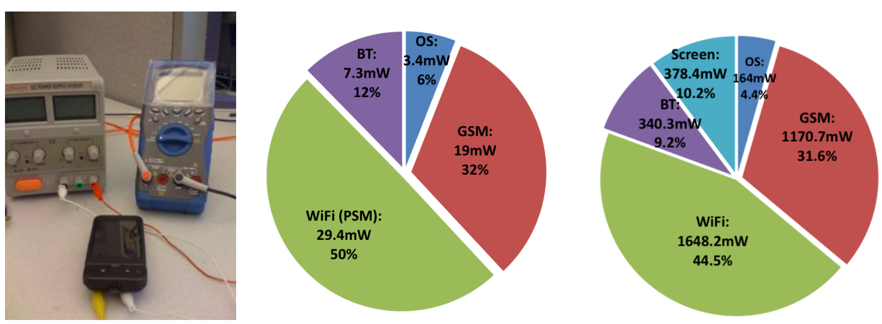

(a) Experimental Setup (b) Radios Idle, Screen Off (c) Radios Active, Screen On

Figure 6: Android G1 Power Consumption Breakdown. (a) our experimental setup to measure the battery power consumption of a smartphone when different network interfaces are turned on and off. (b) Energy Consumption Breakdown in Idle Mode (c) Energy Consumption Breakdown in Active Mode

To have a better understanding of the energy consumption by each components on the latest mobile devices, we measured the energy breakdown on Android G1 phone. Figure 6 shows the power consumption breakdown measured on Android G1 phone in idle and active modes. Although our measurement result shows a great improvements of WiFi energy consumption compared with literature study a few years ago [75,80], WiFi and cellular interfaces are still major energy consumer compared to other components. Particularly, our experiments show that WiFi is very inefficient when no traffic is occurring or when the traffic load is low. This is especially limiting for applications requiring continuous reachability such as VoIP but cannot afford the energy cost of periodic wakeups of WiFi.

Therefore, how to provide users with a reliable, ubiquitous, and yet energy efficient connectivity solution has become a new challenge in today's mobile networks. Our work is focusing on improving the energy efficiency for network communication module. A lot of research efforts have been drawn recently to address this research question. In section 4.3 and 4.2, we will present in detail the related research and our approach to solve this problem.

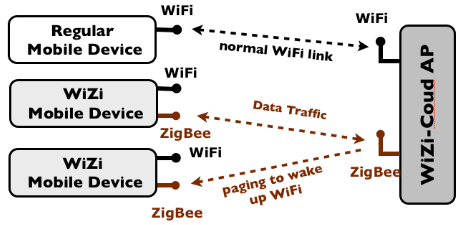

Figure 7: WiZi-Cloud Dual-Radio Solution on Mobile Devices and Access Points. An alternative ultra low-power ZigBee link is established between mobile device and access point for data traffic or signalling.

4.2 Our Approach

With the above constraints bearing in mind, we design and develop WiZi-Cloud, which utilizes ZigBee [81] [1] to establish an alternative ultra low-power wireless link between mobile devices and WiFi access points (See Figure 7). We use this low power link to carry some network traffic when possible, and keep cellular and WiFi off to reduce the energy consumption. In WiZi-Cloud, we design a new hardware which is integrated into both mobile devices and WiFi APs to achieve the energy saving goals with the assistance of the current WiFi infrastructure. Our design of WiZi-Cloud results from the following design decisions:

- **Multi-Radio Solution**: We envision that future mobile phones will be equipped with multiple radios that can connect to the Internet, e.g., current mobile phones already have WiFi, Bluetooth, and GSM. The ZigBee link we propose will co-exist with other network interfaces. Recently a commercial Android mobile device with ZigBee technology was just released [82].
- **Leverage of Existing WiFi Infrastructure:** WiFi, now being commonly deployed, already provides an infrastructure with good coverage. A lot of mobile devices and WiFi APs run on open source software, such as Android and OpenWRT, and can be upgraded. Besides, as hardware becomes commodity, WiFi APs have higher computation capability and better extensibility. For example, all newer WiFi APs have USB interfaces ready. The hardware and software support of the off-the-shelf mobile products make it easy to incrementally deploy some new technology, such as WiZi-Cloud.
- **Choice of ZigBee:** We choose ZigBee because of its zero-time connection establishment, and good radio range (a significant advantage over Bluetooth). ZigBee is also available as a low cost System-on-Chip (SoC) with an integrated low power microcontroller such as in the TI CC2530 [83]. These important features allow the mobile phone to be in sleep mode while the microcontroller handles the wakeup and some of the network functionality with minimum overhead.
- **One Size Does Not Fit All:** Each of these network interfaces has different characteristics in terms of energy consumption, capacity, and coverage. The mobile phone should be able to determine which interface to carry the packets according to its traffic demands and other system conditions. The ZigBee link we prototyped in WiZi-Cloud is an ultra low power link, but has a limited bandwidth compared to WiFi. It is particularly designed for mobile phone applications with moderate traffic demand, or for signalling purpose to wake up the energy consuming radios from sleeping mode.

In this work, we propose the architecture, protocols, and hardware/software implementation of WiZi-Cloud with an emphasis on the following key features.

- **Energy-efficiency**: WiZi-Cloud system is extremely efficient for maintaining connectivity and low rate applications such as VoIP in terms of energy consumption.
- **Leverage of existing HW/SW**: WiZi-Cloud runs on off-the-shelf mobile phones and wireless routers without hardware modifications.
- **Flexibility**: In WiZi-Cloud design, a mobile phone is able to determine the network interface to use according to user-specified policy. WiZi-Cloud provides the mechanism to switch between WiFi and ZigBee interfaces.
- **Seamless**: WiZi-Cloud system and its protocols (e.g., inter-AP handover) is completely transparent to the applications running on the mobile phones and peer entities in the Internet.
- **Coverage**: WiZi-Cloud archives larger coverage than WiFi, and provides two optional extension schemes that can further expand the coverage.

[1] Here we use ZigBee and IEEE 802.15.4 interchangeably.

To the best of our knowledge, this work is the first prototype that integrates ZigBee into commercial cell phones for Internet access. Also, we have conducted comprehensive experiments and measured realistic performance. Our design details, experience, and the evaluation results will certainly benefit and inspire other similar research work in the community.

4.3 Related Works

The literature work on energy consumption on mobile devices can be primarily classified into two categories, 1) optimizing existing network interfaces on mobile devices, by introducing a set of energy saving mechanisms, either designed for general purpose communication or for specific application type, such as VoIP [76,77,84–96]; 2) using alternative low-power wireless link, such as Bluetooth and GSM [75,78–80,97–99], to help improve the energy efficiency. Our prototype considers not only the paging but also the generic data delivery, which provides a more comprehensive energy saving mechanism.

In [78], Shih et al. developed an efficient wake-up mechanism particularly for the VoIP service on PDA-based mobile devices using a special low power control channel between the mobile client and a proxy server. Cell2Notify [75] is another work with the same design goal, but targeting regular cell phones with WiFi capability. In Cell2Notify, WiFi is turned on through the cellular network for the incoming VoIP calls. Both [75] and [78] focus on the paging mechanism that wakes up WiFi for VoIP traffic. And their implementations involves additional hardware such as laptops.

Some other work [99, 100] uses Bluetooth to wake up the WiFi interface. In [100], Agarwal et al.developed a paging scheme assuming each mobile device and the associated AP are connected with a Bluetooth link. Then, WiFi can be turned on via the Bluetooth link. In Blue-Fi [99], the mobile devices predict the availability WiFi connectivity according to the Bluetooth contacts with other nearby Bluetooth devices, and then determine whether to turn on the WiFi. Compared to a Bluetooth link, the Zigbee connection in this paper is significantly superior in terms of handover performance and coverage range. In addition, our system is designed not to wake up the WiFi, but to establish an alternative ZigBee link to carry low rate traffic in a transparent way to the applications.

CoolSpots [80] is a closely related work to this paper. The authors set a Bluetooth link between a mobile device and the associated access point and the traffic can go through either the WiFi or Bluetooth link. CoolSpots focuses on the switching algorithm assuming the bluetooth link has been created by standard. This paper introduces another low-power link using Zigbee which is complimentary with the network interface switching in CoolSpots. In fact, our system can also dispatch packets through different wireless based on specified policy. Furthermore, CoolSpots implements the interface switching by periodically changing the routing rules. Our implementation in this paper supports finer grained control of per packet switch, i.e., the mobile device can determine which network interface to use for each packet.

In addition, there is a large set of work focusing on exploring problematic issues in the current 802.11 protocol. For example, [76,77,86,92,95] improves the current PSM strategy to make WiFi more energy efficient while serving VoIP traffic. Other types of traffic, such as buld data transfer, web browsing, and their energy performance in WiFi networks also have been well studied [84,85,87,89–91,93,94]. In our work, we use ZigBee, which is fundamentally different from WiFi, as a low power alternative radio link to achieve energy efficiency. Our goal is to characterize this new link on mobile devices and quantitatively evaluate its potential and constrains.

4.4 WiZi-Cloud System

4.4.1 System Design and Prototype

The WiZi-Cloud system extends the hardware and network-stack of existing WiFi access points and mobile devices with a set of protocols and mechanisms to support an additional low-power air interface.

Hardware: On the mobile device the ZigBee is integrated as a low cost accessory, in our case interfacing with an Android phone using the serial link. This could be made more compact by using a ZigBee microSD card [101]. We have prototyped a hardware module, *WiZi-Kit*, which integrates TI CC2530, on-board PCB antenna, and connectivity interfaces including UART and FTDI-USB (See Figure 8(a). WiZi- kit can be attached to mobile phones and laptops as a small dongle (See Figure 8(b) and Figure 8(c)). On the AP, we use OpenWRT compatible access points which gives us hundreds of choices from many manufactures [68]. Our current prototype runs on two particular models, Linksys WRT54GL, and Planex Wireless USB router MZK-W04NU (See Figure 9). On WRT54GL, the ZigBee is integrated by soldering four wires on the router board. On the Planex router, the ZigBee dongle can be attached to the USB interface.

Software: The network stack of the access point is extended to maintain connectivity with the mobile devices through the ZigBee interface (e.g., beaconing and paging for ZigBee), as well as to coordinate with peer APs to locate mobile devices. The network stack of the mobile device is extended using a virtual network interface through which all traffic is directed. The network stack maintains connectivity at low energy cost (periodic ZigBee wakeup), and seamlessly switches between the WiFi and ZigBee links using an intra-device handover mechanism depending on the traffic load. It also supports handover

(a) (b) (c)

Figure 8: WiZi-Kit Prototype for Client Devices. (a) WiZi-Kit PCB Design with HTC ExtUSB Interface compatible with Android G1. (b) Android G1 attached with WiZi-Kit. (c) Notebook attached with WiZi-Kit USB Version.

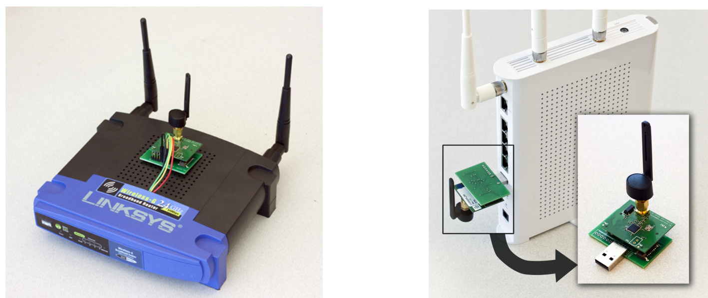

(a) (b)

Figure 9: WiZi-Kit Prototype for Access Points. (a) LinkSys WRT54GL with WiZi-Kit Serial Connection on Board. (b) Planex MZK-W04NU with WiZi-Kit USB version.

across a network of WiZi-Cloud access points as the mobile phone roams around. The network stack extensions are designed to be transparent to the application.

While several previous work considered multi-radio interfaces for energy efficiency in wireless networks, and as we discuss in the related work section, no previous solution achieves our target design objectives in terms of seamless communication, low delay, energy efficiency, and minimal hardware/software modifications.

4.4.2 Software Framework

Figure 10 shows the WiZi-Cloud software framework. We extend the network stack with a *WiZi-Stack*, which is designed to run below the Internet Protocol layer in the TCP/IP model, and above the Link layer. WiZi-Stack consists of three components, WiZi-Cloud service module, WiZi-kit driver and modem logic in ZigBee firmware. Currently, WiZi-stack is implemented as a user space C program.

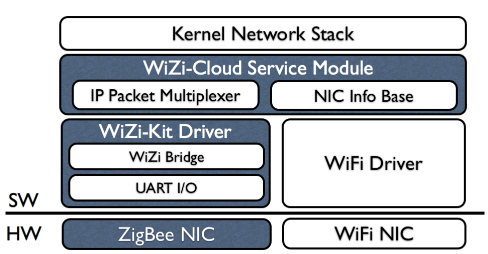

Figure 10: WiZi-Cloud System Framework.

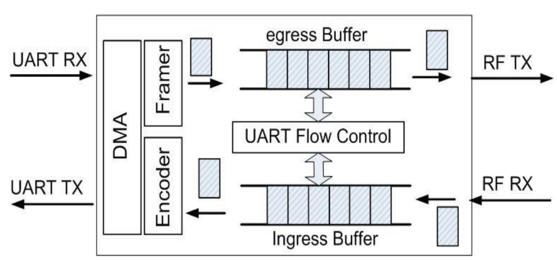

Figure 11: ZigBee Modem Logic

WiZi-Cloud Service Module The main task of this service module is to distinguish the WiZi-Cloud management traffic from generic IP packets and respectively handle them. For regular IP packets, the service module plays the role of a multiplexer passing packets between the kernel network stack and the active radio interface (either WiFi or ZigBee). For WiZi-Cloud management messages, such as registration and paging, the Service Module always forwards them to the ZigBee interface. In addition, WiZi-Cloud Service Module maintains a NIC Information Base (NIB) to track the status of the currently active interface for transmission.

WiZi-Kit Driver WiZi-Kit driver handles the communication between the host and the external ZigBee interface. Since the maximum ZigBee frame payload size is much smaller than IP MTU (In CC2530, the maximum ZigBee frame payload size is 116 byte), WiZi Bridge mainly handles fragmentation for the IP packets. WiZi Bridge chops the IP packets from the Service Layer and get each fragment ready to be transmitted with the ZigBee RF. When receiving an IP packet from the ZigBee interface, WiZi Bridge buffers all the fragments, and forwards the reassembled IP packet to the Service Module.Affiliated with the WiZi Bridge, the UART I/O module is responsible for reliable communication on the UART link between the host and ZigBee device.

ZigBee Modem ZigBee Modem provides the host with read and write operations on the ZigBee link. Figure 11 describes the modem logic implemented on ZigBee. When the host has data to send, it sends the data to the ZigBee interface in the form of UART bit streams, and the data will be transmitted via ZigBee transceiver. Similarly, as the ZigBee transceiver receives a packet from the air, it sends to host through UART link. As we implement the WiZi-Cloud prototype, we learned that it is critical to fully explore the link capacity of both UART and ZigBee radio in order to get good system throughput, which in return has impact on energy efficiency. Therefore, we implement various logic, including CRC, UART flow control, DMA, and etc., to guarantee an efficient and reliable communication link between the host and the ZigBee.

4.5 Performance Evaluation

In this section, we present our preliminary experimental results. To have a comprehensive evaluation of the feasibility of WiZi-Cloud, we evaluate the system from two perspectives, energy consumption and throughput.

Throughput In this experiment, the WiZi-enabled mobile phone is connected to a campus LAN through a WiZi-enabled AP. The destination host is a Linux PC in the same campus LAN. We issue iperf throughput test with a duration of 30 seconds from the mobile phone to the destination host. We measure the throughput for both TCP and UDP, with varying UDP payload size and TCP MSS. Figure 12 shows the throughput results. Each data point is the iperf reported goodput, averaged over 10 runs.

- *UDP Throughput*: When the UDP payload size is small, the WiZi-Cloud protocol header incurs a large overhead yielding a low throughput. As the payload size increases, the overhead is amortized and the throughput quickly increases. Beyond 500 byte, the curve shows limited growth because the whole data flow along the UART and ZigBee radio link is efficiently pipelined. The peak UDP throughput is 70.4Kbps with 1400 Byte payload. With UART link header overhead accounted, this translates to UART link throughput 83Kbps, which is close to the G1 UART link speed limit (115Kbps). This shows that our prototype implementation is efficient.
- *TCP Throughput*: Similar to UDP, the TCP MSS also has impact on the ZigBee link utilization. Therefore, the TCP throughput shows similar growing curve as the MSS increases. However, the TCP throughput starts dropping as the MSS keeps increasing. With larger MSS, more IP fragments are generated while transmitting them on ZigBee link. Due to the limited processing power of ZigBee chipset and slow UART link, IP fragments may be lost when the buffer is overrun. If one IP packet fragment is lost, all the rest of the fragments will be of no use. Thus, the TCP MSS becomes a trade off between better channel utilization and the risk of wasting bandwidth. In our experiment, the optimal TCP MSS is 450 Byte, achieving 60.2Kbps throughput.

<table>
<tr><th></th><th>GSM</th><th>WiFi</th><th>Bluetooth</th><th>ZigBee</th><th>OS</th><th>Screen</th></tr>
<tr><td>RF Idle</td><td>19.04</td><td>29.42</td><td>7.32</td><td>2.57</td><td rowspan="2">3.54</td><td rowspan="2">378.144</td></tr>
<tr><td>RF Active</td><td>1170.71</td><td>1648.2</td><td>340.3</td><td>94.5</td></tr>
</table>

Table 1: Breakdown Energy Consumption on Android Phones in mW.

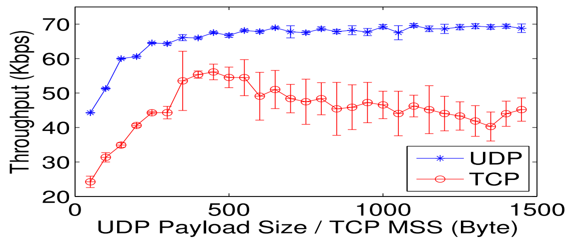

Figure 12: WiZi TCP/UDP Goodput with varying TCP MSS and UDP Payload Size.

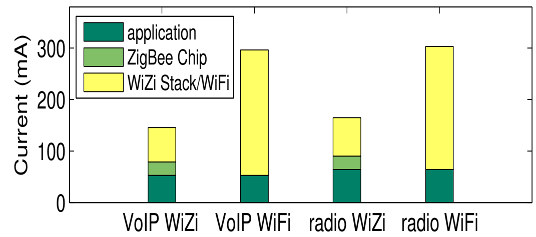

Figure 13: Total Energy Consumption of SipDroid VoIP Application in Active and Standby Mode. WiZi vs. WiFi.

Energy Consumption The energy consumption is one of the most important metrics in our experiments. First, we show the breakdown of energy consumption measured with Android G1 in Table 1. To measure the phone energy consumption, we power the phone with an external power generator (4.1V), and connects the Agilent U1252A multimeter in series. The multimeter logs the instantaneous current value every 5ms. The result shows that ZigBee in idle mode achieves more than factor of 11 improvement in energy consumption in comparison with WiFi in Power Saving Mode, and a factor of 7 in comparison with GSM.

However, energy usage of the radio interface cannot tell the whole story. Due to the low data rate and limited computation capability of ZigBee chipset, it may not be suitable for all applications and it is important to study how ZigBee would impact the overall system energy usage.

We tested a VoIP application called SipDroid with two popular codecs, GSM 13Kbps and Speex 11Kbps. The voice is clear, however sipdroid does not report any statistical data indicating the call quality. We capture the SipDroid traffic, and use iperf to emulate the VoIP traffic by generating two-way UDP flow, with proper packet size and packet rate.

Figure 13 shows the energy consumption by SipDroid in active mode (during VoIP call) and standby mode, with WiFi and ZigBee, respectively. Each bar consist of three components: 1) the base energy usage, including the energy consumed by OS, speaker and application; 2) the energy consumed by the WiFi or the whole WiZi-Stack; 3) the energy consumed by the external ZigBee hardware (none in WiFi case). In this type of applications, packets come at a fast pace, which prevents both WiFi and WiZi from entering the power save mode. This results in a high WiFi energy consumption, of around 250mA in both applications. In contrast, ZigBee consumes only around 27mA even in active mode. In standby mode, WiZi system shows even higher energy efficiency because the energy usage by the ZigBee hardware and WiZi-Stack is very negligible. The phone standby time with VoIP software is extended by three times.

4.6 Future Work

Our preliminary results show the promising improvement of energy consumption, and good throughput performance, which make WiZi-Cloud suitable for a potential set of applications. We propose the following steps to complete this research work.

- *Evaluation of Other Mobile Applications*: Since different mobile applications feature different traffic load, level of user interactivity, etc., WiZi-Cloud solution may not be suitable for all applications. We plan to evaluate a set of mobile applications to verify the suitability of WiZi-Cloud on mobile platforms.
- *Coverage Performance*: ZigBee transmits in ultra low power. In order to provide a ubiquitous and reliable communication in urban area, the coverage is the key factor. we plan to carry out an extensive experiment to evaluate the coverage performance of WiZi-Cloud. In addition, we plan to explore several potential mechanisms to extend the coverage.

5 BaPu - Practical Bunching of Access Point Uplinks

5.1 Motivation

Nowadays, the mobile devices are equipped with high-resolution cameras and a variety of sensors, and are quickly becoming the primary device to generate personal multimedia content. Both the quality and quantity of the User Generated Content (UGC) continuously grow. This naturally leads to users' ever increasing demand of sharing these high volume of UGC with others in an instant way, or backing them up in the "Cloud Storage". To obtain a satisfactory user experience, users need sufficient uplink bandwidth to do the fast bulk data transfer to the Internet. However, today's ISPs generally offer highly throttled uplink bandwidth of around 1 to 3 Mbps. As a result, the instant sharing of HD content or the fast data backup in

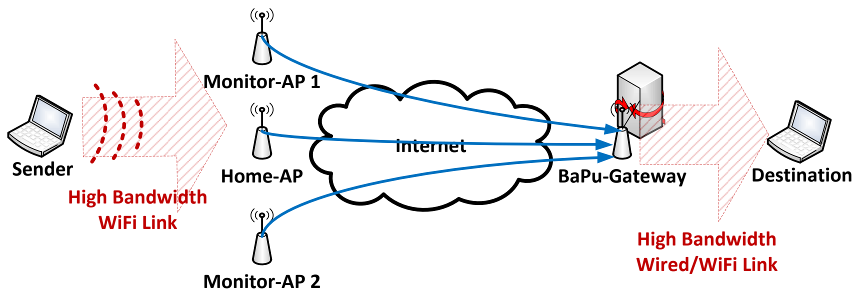

Figure 14: BaPu System Architecture

the "Cloud" is generally infeasible in today's residential broadband. For example, with 3Mbps uplink, it takes about 1.2 hours to upload a 30 minute 1080p video clip (8 Mbps bitrate). Even though the consumer device manufactures and cloud service providers have high demand of better uplink speed to enhance the user experience, after all it will place a big burden on ISP to upgrade their networks to provide higher uplink speed. We believe this is less likely to have a much faster uplink in the near future without any ISP-friendly business model is invented.

In this work, we examine the mechanisms of aggregating multiple broadband uplinks, with the assistance of the WiFi infrastructure in the same neighbourhood. Our core idea primarily relies on the following observations,

- **Asymmetric WiFi bandwidth and broadband uplink**: In contrast to access points backhaul, WiFi wireless bandwidth is much larger. 802.11n can support up to 600Mbps data rate. With sufficiently high WiFi bandwidth, it is beneficial to aggregate the proximate wired broadband connections using WiFi communication with multiple APs.
- **WiFi densely deployed in residential area**: The density of WiFi APs is very high in residential areas. This number exceeds 10 APs per points already in 2005 [102]. Our recent Wardriving measurements, conducted in late 2011, indicate an average of 17 APs in the city of Boston [26]. Such widely present WiFi APs, to some extent, justifies the feasibility of backhaul aggregation through WiFi.
- **Mostly idle broadband:** As we discussed in section 3.2, we observe that the backhaul traffic load is usually very low, even at peak hours. This makes traffic multiplexing a viable approach for scaling backhaul capacity.

5.2 Our Approach

5.2.1 Overview of BaPu

The primary design goal of BaPu is to aggregate the limited backhaul uplink bandwidth behind each WiFi AP by having the sender communicate with multiple WiFi APs simultaneously. In our design, we take advantage of the broadcast nature of wireless communication, and the high WiFi bandwidth. Figure 14 shows the BaPu system architecture. As shown in the architecture, the *sender* and its *home-AP* reside in the same WLAN. Other proximate APs residing in the same neighbourhood are denoted as *monitor-AP*. Both home-APs and monitor-AP are called BaPu-AP. Each BaPu-AP is configured as both AP mode and monitor mode. Having one physical WiFi interface running in multiple modes are generally supported by today's WiFi drivers, such as Atheros. In our design, each BaPu-AP may act as a home-AP for all the clients inside the same WLAN, or a monitor-AP to carry the traffic for the clients from other WLANs. As sender uploads data to the remote destination through its home-AP, the wireless link can provide a very high wireless bandwidth, generally tens of Mbps. Meanwhile, the monitor-APs overhear the traffic. The home-AP and monitor-APs collaborate with each other and each AP uploads a share of the traffic to the remote destination. By aggregating multiple backhaul uplink throughput, the end to end throughput between the sender and the destination can be boosted. In BaPu, the single data flow sourcing from the sender is split into multiple flows and forwarded to the destination. Thus, we design a BaPu-Gateway, which runs in front of the destination node. BaPu-Gateway is responsible for reassembling the separate data flows from multiple BaPu-APs, and forward the restored single data flow to the actual destination. In this way, the bandwidth aggregation is totally transparent to both sender and destination. The big advantage of such transparency is that no modification is required from the existing network stacks and network applications running on the sender and destination. In our design, the BaPu-Gateway is also a pure software solution that can be run on the WiFi AP of the destination, or a gateway server on the edge of some online service's server clusters, e.g. Youtube and DropBox. However, the transparency design and the fact that single data flow is split into multiple flows raise many technical challenges which makes the efficient multi-flow traffic forwarding non-trivial.

5.2.2 Mechanisms

To achieve the design goal of transparent and efficient multi-flow backhual uplink aggregation, we design a complete suite of mechanisms.

Transparency and Tunnel Forwarding

One of our design goals is to make BaPu a transparent solution for the sender and the destination, and no modification is required from them. Therefore, pure software upgrade on the WiFi APs and gateway servers can be commonly supported on today's commodity hardware, and the incremental deployment of BaPu becomes much easier. As BaPu-APs identify a new session, either TCP or UDP, established between the sender and a certain destination, the BaPu-AP first establishes a TCP tunnel connection to the BaPu-Gateway. Through this TCP tunnel, each BaPu-AP exchanges control information with the BaPu-Gateway and forwards the data traffic to it. The choice of TCP tunnel is determined by the following two factors: 1) Different BaPu APs generally have different IP addresses. As the monitor-AP receives IP packets from the sender, it cannot simply forward the unchanged IP packet to the destination, because such IP packet has the source IP address of home-AP and most ISPs block IP spoofing traffic. Therefore, the traffic must be tunnelled to the destination. 2) TCP is designed to be Internet-friendly. The congestion control and flow control mechanisms of TCP can prevent the Internet from being overloaded as we aggregate the backhual uplink.

Scheduling and Load Balancing

BaPu-APs may have varying broadband connections for the following reasons: 1) BaPu-APs run on different ISPs, with different broadband subscription plan, which results in different maximum downlink/uplink capacity; 2) The traffic load happening inside one WLAN may constantly change at different time of day. In order not to affect the normal network usage, the amount of idle bandwidth that a certain BaPu-AP may contribute is a dynamic factor; 3) the network latency between the BaPu-AP and destination may vary as well due to similar reasons we previously discussed. Therefore, it is critical to schedule the traffic share among APs in order to ensure the efficient and fair bandwidth aggregation. In BaPu, we adopt a centralized scheduling and load balancing mechanism at BaPu-Gateway. As shown in Figure 15(a), as BaPu-AP receives packet streams from the sender, it identifies each IP packet with IP ID (<IP ID, TCP seq> tuple for TCP flow). BaPu-AP first sends a reception report for each packet to the BaPu-Gateway. BaPu-Gateway, therefore, has the complete knowledge of which BaPu-AP receives which IP packet. Along with the reception report, each BaPu-AP also periodically obtains the TCP stack statistics info from the Linux kernel stack and sends the local idle bandwidth estimate to the BaPu-Gateway. Based on the full reception report and the idle bandwidth estimate, BaPu-Gateway determines the packet schedules, and notifies all APs. In this way, we can efficiently utilize the idle bandwidth of each BaPu-AP, as well as well adapts to the changing traffic load on each AP.

Special Challenges of TCP and BaPu-Pro

As we design and implement BaPu, we observe that the scheduling and load balancing mechanisms work well for UDP, and efficiently aggregate 95% of total idle bandwidth. However, the throughput downgrades a lot for TCP, and achieves only around 50% of total bandwidth with 7 BaPu-APs. Our in-depth analysis reveals that some TCP inherent nature raises some challenges specific to TCP, and limits the aggregated throughput performance.

TCP is designed to ensure successful and in-order data delivery between the sender and the receiver, relying on two major mechanisms, *flow control* and *congestion control*. The former one prevents the sender from overrunning the receiver buffer. The latter one prevents the sender from overrunning the network path between the sender and the receiver. In TCP, each packet is identified with a TCP sequence number and must be acknowledged by the receiver to indicate the proper delivery. The sender maintains a dynamic CWND (Congestion Window) as the communication goes on, which indicates the maximum number of packets on the fly. If some packet is not ACK'ed before RTO (Retransmission Timeout) occurs, the sender issues retransmission. If the receiver receives some out-of-order sequence number, this generally implies the missing sequence number is lost or delayed due to congested network. In this case, the receiver sends a DUPACK (Duplicate ACK) back to sender to notify such missing sequence occurrence. By default, the sender issues Fast_Retransmission upon receiving 3 DUPACKs. In either of these two cases, the sender generally reduces the CWND accordingly to slow down the sending rate and adapt to the congested network path or slow receiver.

The traditional TCP was designed based on the fact that the out-of-order sequence number is generally a good indication of packet being lost or the network path is experiencing congestion during that moment. However, such assumption no longer holds in BaPu. In BaPu, the packets belonging to the same TCP flow are intentionally routed through multiple APs, which generally have diverse broadband connections with different capacity, network latency, traffic load, etc. This results in very serious out-of-order packet arrival at the BaPu-Gateway. As BaPu-Gateway forwards the out-of-order packets to the destination, the sender and destination falls on the TCP mechanisms to limits the CWND at a very low value. Thus, the end to end throughput is highly limited.

To tackle the TCP specific challenges, we propose a Proactive-ACK mechanism. To differentiate these two versions, we denote them as BaPu-Basic and BaPu-Pro, respectively. As shown in Figure 15(b), as BaPu-Gateway receives packet reception report, it maintains a lookup table for all the reported TCP sequence number. Since the tunnel connection between BaPu-AP and BaPu-Gateway is TCP-based, the IP packet with the reported sequence for sure will be delivered ultimately. Therefore, as long as the BaPu-Gateway identifies a continuous range of reported sequence numbers, it pro-actively sends spoofed TCP ACKs back to sender, without receiving the actual IP packet. Meanwhile, BaPu-Gateway buffers the out-of-order sequence numbers and only forwards the continuous sequence to the destination. In this way, we can have sender CWND grows in a healthy way, and keep the destination not ware of such out-of-order packets. Our experiment results presented in Section 5.4 show the efficacy of BaPu-Pro while serving TCP traffic.

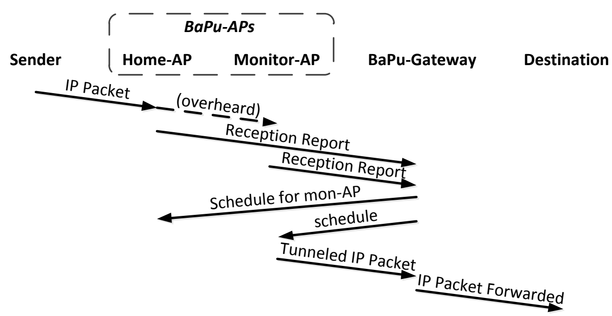

(a) BaPu-Basic Protocol Flow: Scheduling and Load Balancing

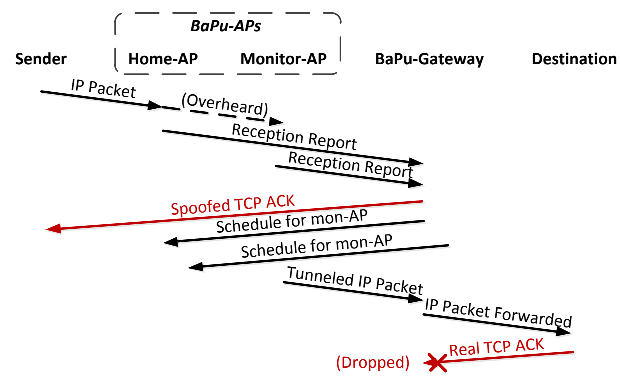

(b) BaPu-Pro Protocol Flow: Proactive-ACK

Figure 15: BaPu Protocol Flow

5.3 Related Works

Exploiting multiple radio links to improve devices throughput has been explored from several perspectives including traffic aggregation [54], mutipath forwarding [52], mitigation of wireless losses [103]. In addition to systems that rely on multiple radio interfaces [79], many solutions and algorithms were proposed for a single radio interface that carefully switches across multiple access points while providing the upper layers of the network stack a transparent access through a virtual interface [49, 54, 104, 105]. Solutions to overcome the limited APs backhaul through aggregation using a virtualized radio interface include the initial Virtual-WiFi system where two TCP connection might be services by through two different APs [105], FatVAP that achieves fast switching a smart AP selection [54], ARBOR that add security support [49], Fair WLAN that provides fairness [47]. Many of these systems require techniques for fast switching across access points to reduce the impact on TCP performance in terms of delay and packet loss as proposed in Juggler [50], and WiSwitcher [51]. The closest approach to our work is the Link-alike system where access points coordinate to opportunistically schedule the traffic over their backhaul links [52]. Our approach differs from previous work in requiring that the client devices remain unchanged and requires the transparent support of protocols such as TCP.

Being completely transparent to the clients and constraining each link AP-Destination flow to be TCP-friendly makes efficient multipath transport, a key component of our system. There has been a significant amount of work in this area for quite some time from various perspectives that are fairly different from our setup. Previous work identified the issues with differential delay, and rates and mostly focussed on providing reliability, flows balancing, and maintaining fairness. Proposed solutions, require the modification of the client devices network stacks, and usually do not aim at increasing capacity through simultaneous use of all paths. For example, the IETF has two standards on transport protocols using multipath. The Stream Control Transmission Protocol (SCTP) was primarily designed for multi-homed devices that require fail-over support [106], the more recent Multi-Path TCP (MPTCP) is a TCP extension that aims at enabling nodes to efficiently communicate utilizing multiple parallel paths [107]. In recent work, [108] proposed a congestion control mechanism for MPTCP with the objective of reliability and fairness. Other transport protocols that require the modification of the client devices, include pTCP [109] an end to end transport protocol that achieves bandwidth aggregation on multi-homed mobile hosts, RCP [110] a Receiver Centric Protocol that aggregates heterogeneous interfaces traffic and carries link specific congestion control, R-MTP [111] balances and coordinates the traffic over wireless links with varying characteristics, Horizon [112] uses back- pressure technique to balance the traffic across forwarding paths. Beyond mobile communication environments, multipath TCP features have also been finding applications in various networking environments such as data [113,114]. The distinguishing element of BaPu is that it aims at transparently supporting unmodified client devices and TCP/IP stacks while efficiently aggregating the APs backhauls.

5.4 Performance Evaluation

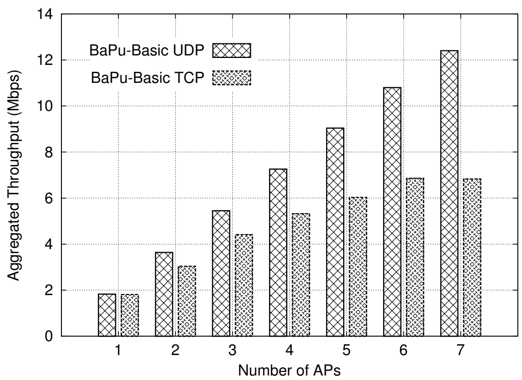

(a)

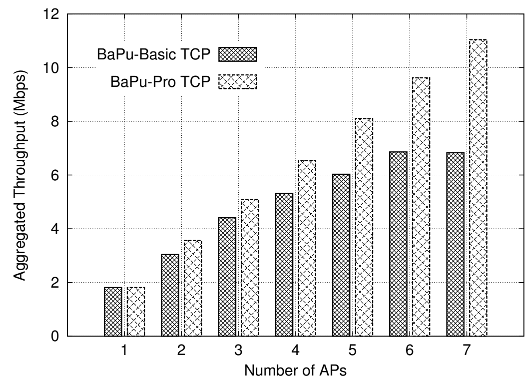

(b)

Figure 16: BaPu Throughput Experiment: (a) BaPu-Basic Aggregated Throughput Performance for UDP and TCP. (B) TCP Aggregated Throughput Performance, BaPu-Basic vs. BaPu-Pro.

We have implemented our prototype BaPu system with commodity hardware. Our testbed consists of one laptop client, 7 BaPu-APs, 1 BaPu-Gateway and 1 destination node. Each AP is a Buffalo WZR-HP-G300NH 802.11n wireless router. We have reflashed the APs with OpenWRT firmware, running Linux with kernel 2.6.32 and the ath9k Wi-Fi driver. In our experiments, 1 BaPu-AP acts as home-AP which the client is always associated with. The other 6 BaPu-APs act as neighbouring monitor-APs to capture the traffic in monitor mode. In our implementation, BaPu-Gateway runs on a Linux PC, and the destination node is running behind the Gateway as an in-LAN node. The legacy client and the destination node are both laptops with a 802.11n WiFi card, running Ubuntu Linux. All BaPu-APs and BaPu-Gateway are inter-connected. To emulate the traffic shaping in residential broadband, we install a Linux PC between the BaPu-APs and the BaPu-Gateway. We use *tc* to shape the downlink/uplink bandwidth (*htb*) and latency (*netem*). The bandwidth is throttled to be 20Mbps and 2Mbps for downlink and uplink, respectively. The round-trip latency is shaped to be 32ms. All parameters are obtained by our measurement on the Open Infrastructure testbed.

We measure the throughput performance by issuing long-term iperf traffic flows between the client and destination. Each iperf tests last for 1 hour. Figure 16(a) shows the throughput performance of BaPu-Basic for UDP and TCP. As the number of APs increases, the UDP aggregated throughput increases proportionally. We achieve up to 12.4Mbps, which translates to 97.33% of theoretical maximum aggregated throughput with all the BaPu protocols overhead taken into account. This also indicates the efficiency of our implementation. In the contrast, the TCP aggregated throughput growth slows down as the

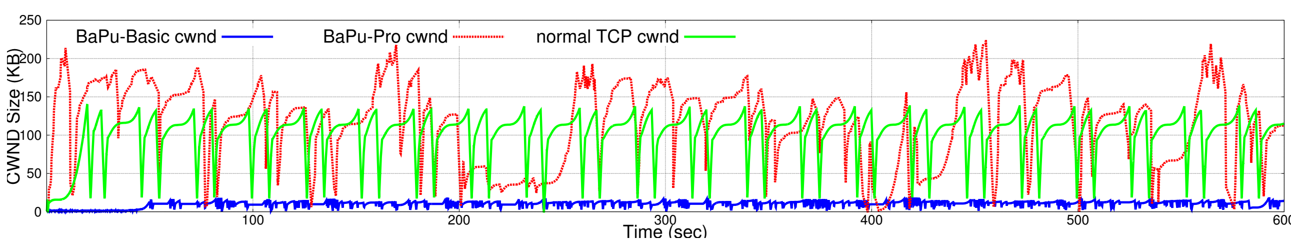

Figure 17: Comparison of TCP CWND (Congestion Window) growth during 600 second iperf test. CWND value remains small in BaPu-Basic, resulting in 5Mbps aggregated throughput. In BaPu-Pro, CWND can grow to much larger value, and achieves 11Mbps. BaPu-Pro CWND growth follows a very similar pattern as normal TCP flow with throughput throttled to be 11Mbps.

AP number increases. We achieve only 51% of theoretical maximum aggregated throughput. Figure 16(b) shows a significant improvement of TCP throughput by BaPu-Pro. With 7 APs, we achieve 11.04Mbps aggregated throughput, 88% of the upper bound value.

To further justify our design choice of the Proactive-ACK mechanisms in BaPu-Pro, we log the TCP stack state variable at the sender client by reading the TCP_INFO from Linux kernel with getsockopt(). As shown in Figure 17, in BaPu-Basic, the CWND size remains very small as the TCP session goes on. In the contrast, BaPu-Pro allows the CWND grow to very high value, which boosts up the throughput. For reference purpose, we also run a normal TCP session with a throttled bandwidth 11Mbps (similar to the BaPu-Pro resulting throughput). The CWND growth patterns for BaPu-Pro and normal TCP is very close, which also implies our BaPu-Pro design and implementation can efficiently and transparently aggregate multiple slow uplinks.

5.5 Future Work

Even though our preliminary experiments show promising results, there are many dynamic factors in real residential network and they may have impact on BaPu performance. We plan to carry out an extensive set of experiments to evaluate the BaPu performance.

- We plan to study how BaPu adapts to the changing traffic load in the broadband connections.
- Throughput does not tell all the story. Sufficient throughput is good enough for large data transfer, but may not be suitable for real-time applications, which has high delay sensitivity. We plan to study how BaPu performs for real-time applications, such as HD online streaming.
- With multiple WiFi APs collaborate together, we would like to study the potential impact of the wireless diversity in the lossy urban WiFi environment.
- The actual potential performance in real urban WiFi environment may be determined by a set of factors, including the density of WiFi AP, normal traffic load of WiFi APs, packet reception rate while overhearing other WLANs, etc. We plan to carry out a data-trace driven evaluation on our Open Infrastructure testbed in order to get a more accurate estimate of the potential performance of BaPu.

6 Schedule

The following table is the proposed timeline to complete the research:

To-do tasks	Completion Date
WiZi-Cloud performance evaluation on major mobile applications	September 2012
WiZi-Cloud coverage evaluation	September-October 2012
BaPu Comprehensive Evaluation	October-November 2012
Urban WiFi data trance analysis to evaluate the feasibility of BaPu system	November 2012
Ph.D. Dissertation Defense	January 2013

References

[1] Smart phones overtake client PCs in 2011. *Canalys Press Release*, 2012.

[2] Smartphone Market Forecast: Sales Will Exceed 1.5 Billion Units A Year By 2016. *Business Insider Research Report*, 2012.

[3] Kristin Burnham. ï£ijFacebook Users by the Numbers. 2012.

[4] Dropbox Press articles. *Dropbox*, 2012.

[5] It's YouTube's 7th birthday.. *Youtube Blog*, 2012.

[6] IDC: More Mobile Internet Users Than Wireline Users in the U.S. by 2015. *IDC Press Release*, 2012.

[7] Cisco. Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016. *Cisco White Paper*, 2012.

[8] Apple Genius Bar: iPhones' 30% Call Drop Is Normal in New York. `http://gizmodo.com/5370493/apple-genius-bar-iphones-30-call-drop-is-normal-in-new-york`.

[9] iPhones Overload AT&T's Network, Angering Customers - NYTimes. *New York Times*, 2012. `http://www.nytimes.com/2009/09/03/technology/companies/03att.html`.

[10] iPhone vs. Pre: Satisfaction Bakeoff.
http://brainstormtech.blogs.fortune.cnn.com/2009/08/14/iphone-vs-pre-satisfaction-bakeoff/.

[11] Moray Rumney. Identifying technology to deliver the next 100x capacity growth in wireless. *The 3rd LTE World Summit*, 2008.

[12] Network Investment Infographic 2012. *AT&T Intellectual Property*, 2012.

[13] Cisco. Evolution of the Mobile Network. *Cisco White Paper*, 2010.

[14] Sarah Perez. Mobile data traffic surge: 40 exabytes by 2014. *Read Write web blog*, 2010.
http://www.readwriteweb.com/archives/mobile_data_traffic_surge_40_exabytes_by_2014.php.

[15] A Balasubramanian, R Mahajan, and A Venkataramani. Augmenting Mobile 3G Using WiFi. *Proc. of MobiSys*, 2010.

[16] Kyunghan Lee, Joohyun Lee, Yung Yi, Injong Rhee, and Song Chong. Mobile data offloading: how much can WiFi deliver? In *Co-NEXT '10: Proceedings of the 6th International COnference*, 2010.

[17] Scherzer Shimon. Femtocell vs. Wi-Fi for Data Offloading & Indoor Coverage. *Wireless Week*, 2010.
http://www.wirelessweek.com/Articles/2010/02/Technology-Data-Offloading-Indoor-Coverage-Femtocell-Wi-Fi.

[18] Wireless Data Offloading Solutions Overview. *Greenpacket*.
http://www.greenpacket.com/Mobile_Data_Offload_Solutions_overview.html.

[19] More Than 12 Million AT&T, Starbucks Customers to Get Free Wi-Fi Access for a Rich In-Store Experience. *AT&T Press Release*, 2008.
http://www.att.com/gen/press-room?pid=4800&cdvn=news&newsarticleid=25152.

[20] FON.
http://corp.fon.com/us/.

[21] Small cell market to hit 3 million units in 2016, driven by urban area capacity upgrades. *Infonetics Research Report*, 2012.

[22] Mobile operators ramp WiFi deployments for data offload, seek to integrate WiFi. *Infonetics Research Report*, 2012.

[23] Cisco. The Future of Hotspots: Making Wi-Fi as Secure and Easy to Use as Cellular. *Cisco White Paper*, 2012.

[24] Cisco. Architecture for Mobile Data Offload over Wi-Fi Access Networks. *Cisco White Paper*, 2012.

[25] The Future of WiFi Connection: 802.11u. *FON Blog*.
http://blog.fon.com/en/archivo/technology/future-of-wifi.html.

[26] Open Infrastructure.
http://www.ccs.neu.edu/home/noubir/projects/openinfrastructure.

[27] Wardriving - Wikipedia.
http://en.wikipedia.org/wiki/Wardriving.

[28] IEEE 802.11u. *Wikipedia*.
http://en.wikipedia.org/wiki/IEEE_802.11u.

[29] Hotspot 2.0 and the Next Generation Hotspot. *Certified Wireless Network Professional*.
http://www.cwnp.com/cwnp_wifi_blog/hotspot-2-0-and-the-next-generation-hotspot/.

[30] IEEE 802.11ac. *Wikipedia*.
http://http://en.wikipedia.org/wiki/IEEE_802.11ac.

[31] Aldo Cassola, Tao Jin, Harsh Kumar, Guevara Noubir, and Kamal Sharma. SNEAP: A Social Network-Enabled EAP Method: No More Open Hotspots. *NSDI 2011 Demo*, 2011.

[32] Vytautas Valancius, Cristian Lumezanu, Nick Feamster, Ramesh Johari, and Vijay V Vazirani. How Many Tiers? Pricing in the Internet Transit Market. *Proceedings of SIGCOMM*, 2011.

[33] R Di Taranto and C Rosenberg. Throughput-based incentives for residential femtocells. *Personal Indoor and Mobile Radio Communications (PIMRC), 2011 IEEE 22nd International Symposium on*, 2011.

[34] Kideok Cho, Hakyung Jung, Munyoung Lee, Diko Ko, T.T Kwon, and Yanghee Choi. How can an ISP merge with a CDN? *Communications Magazine, IEEE*, 2011.

[35] Nikhil Gopinath Shetty Jean Walrand. Design of network architectures: role of game theory and economics. *Design of network architectures: role of game theory and economics*, 2010.

[36] V Misra, S Ioannidis, A Chaintreau, and L Massoulié. Incentivizing peer-assisted services: A fluid Shapley value approach. *ACM SIGMETRICS Performance Evaluation Review*, 2010.

[37] G Biczók, L Toka, A Vidacs, and T A Trinh. On incentives in global wireless communities. *Proceedings of the 1st ACM workshop on User-provided networking: challenges and opportunities*, 2009.

[38] M. El Dick, E. Pacitti, and B. Kemme. A highly robust P2P-CDN under large-scale and dynamic participation. *Advances in P2P Systems, 2009. AP2PS'09. First International Conference on*, 2009.

[39] K Wu, H Tan, Y Liu, J Zhang, Q Zhang, and L Ni. Side channel: bits over interference. *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, 2010.

[40] X Yue, CF Wong, and SHG Chan. A distributed channel assignment algorithm for uncoordinated WLANs. *Consumer Communications and Networking Conference (CCNC), 2010 7th IEEE*, 2010.

[41] R Murty, J Padhye, A Wolman, and M Welsh. Dyson: an architecture for extensible wireless LANs. *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, 2010.

[42] I Broustis, K Papagiannaki, SV Krishnamurthy, M Faloutsos, and VP Mhatre. Measurement-Driven Guidelines for 802.11 WLAN Design. *Networking, IEEE/ACM Transactions on*, 2010.

[43] MK Panda and A Kumar. Modeling multi-cell IEEE 802.11 WLANs with application to channel assignment. *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2009. WiOPT 2009. 7th International Symposium on*, 2009.

[44] V Shrivastava, N Ahmed, S Rayanchu, S Banerjee, S Keshav, K Papagiannaki, and A Mishra. Centaur: Realizing the full potential of centralized wlans through a hybrid data path. *Proc. of MobiCom*, 2009.

[45] R Murty, J Padhye, R Chandra, A Wolman, and B Zill. Designing high performance enterprise wi-fi networks. *proc. of USENIX*, 2008.

[46] Hamed Soroush, Peter Gilbert, Nilanjan Banerjee, Mark D Corner, Brian N Levine, and Landon Cox. Spider: improving mobile networking with concurrent wi-fi connections. In *SIGCOMM '11: Proceedings of the ACM SIGCOMM 2011 conference on SIGCOMM*, 2011.

[47] Domenico Giustiniano, Eduard Goma, Alberto Lopez Toledo, Ian Dangerfield, Julian Morillo, and Pablo Rodriguez. Fair WLAN backhaul aggregation. In *MobiCom '10: Proceedings of the sixteenth annual international conference on Mobile computing and networking*, 2010.

[48] Fahad R Dogar, Peter Steenkiste, and Konstantina Papagiannaki. Catnap: exploiting high bandwidth wireless interfaces to save energy for mobile devices. In *MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services*, 2010.

[49] Xinyu Xing, Shivakant Mishra, and Xue Liu. ARBOR: hang together rather than hang separately in 802.11 wifi networks. In *INFOCOM'10: Proceedings of the 29th conference on Information communications*, 2010.

[50] A.J Nicholson, S Wolchok, and B.D Noble. Juggler: Virtual Networks for Fun and Profit. *Mobile Computing, IEEE Transactions on*, 2010.

[51] Domenico Giustiniano, Eduard Goma, Alberto Lopez, and Pablo Rodriguez. WiSwitcher: an efficient client for managing multiple APs. In *PRESTO '09: Proceedings of the 2nd ACM SIGCOMM workshop on Programmable routers for extensible services of tomorrow*, 2009.

[52] Szymon Jakubczak, David G Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Link-alike: using wireless to share network resources in a neighborhood. *SIGMOBILE Mobile Computing and Communications Review*, 2009.

[53] Paramvir Bahl, Ranveer Chandra, Patrick P. C. Lee, Vishal Misra, Jitendra Padhye, Dan Rubenstein, and Yan Yu. Opportunistic use of client repeaters to improve performance of WLANs. In *CoNEXT '08: Proceedings of the 2008 ACM CoNEXT Conference*, 2008.

[54] S Kandula, KCJ Lin, T Badirkhanli, D Katabi, MIT NTU, and M MIT. FatVAP: aggregating AP backhaul capacity to maximize throughput. *Proc. of the USENIX NSDI*, 2008.

[55] D Cho, A Wolman, and S Saroiu. MAUI: Making Smartphones Last Longer with Code Offload. *Proc. of MobiSys*, 2010.

[56] M Marcon, B Viswanath, M Cha, and KP Gummadi. Sharing Social Content from Home: A Measurement-driven Feasibility Study. *proceedings of NOSSDAV*, 2011.

[57] J Tong, K Xu, and R Pi. A new Web Service Structure of Combining P2P and CDN Technologies. *Web Society (SWS), 2010 IEEE 2nd Symposium on*, 2010.

[58] H Yin, X Liu, T Zhan, V Sekar, F Qiu, C Lin, H Zhang, and B Li. Design and deployment of a hybrid CDN-P2P system for live video streaming: experiences with LiveSky. *Proceedings of the seventeen ACM international conference on Multimedia*, 2009.

[59] V Valancius, N Laoutaris, L Massoulié, C Diot, and P Rodriguez. Greening the internet with nano data centers. *Proc. of CoNEXT*, 2009.

[60] Watch Global, Cache Local: YouTube Network Traffic at a Campus Network - Measurements and Implications. 2011.

[61] S Kang and H Yin. A hybrid CDN-P2P system for Video-on-Demand. *2010 Second International Conference on Future Networks*, 2010.

[62] D Xu, S Suresh, and K Catherine. Analysis of a CDN-P2P Hybrid Architecture for Cost-Effective Streaming Distribution," ACM/Springer Multimedia Systems. *Journal*, 2008.

[63] C Huang, A Wang, J Li, and KW Ross. Understanding hybrid CDN-P2P: why limelight needs its own Red Swoosh. *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, 2008.

[64] V Aggarwal, A Feldmann, and C Scheideler. Can ISPs and P2P users cooperate for improved performance? *ACM SIGCOMM Computer Communication Review*, 2007.

[65] D Xu, SS Kulkarni, C Rosenberg, and HK Chai. Analysis of a CDN–P2P hybrid architecture for cost-effective streaming media distribution. *Multimedia Systems*, 2006.

[66] E. Turrini and F. Panzieri. Using P2P techniques for content distribution internetworking: a research proposal. *Peer-to-Peer Computing, 2002.(P2P 2002). Proceedings. Second International Conference on*, 2002.

[67] OpenWRT. OpenWRT – Wireless Freedom, 2012.
`http://openwrt.org/`.

[68] OpenWRT – Table of Hardwares.
`http://wiki.openwrt.org/toh/start`.

[69] Buffalo WZR-HP-G300NH.
`http://www.buffalotech.com/products/wireless/single-band-wireless-routers`.

[70] The Multiscreen Revolution. *Adobe Featured Blogs*, February 22, 2011.
`http://blogs.adobe.com/conversations/2011/02/the-multiscreen-revolution.html`.

[71] Why Your Smartphone Battery Sucks. *PCWorld*.
`http://www.pcworld.com/article/228189/why_your_smartphone_battery_sucks.html`.

[72] Report Breaks Down Top Smartphone Complaints. *PCMAG.com*.
`http://www.pcmag.com/article2/0,2817,2406804,00.asp`.

[73] Vladimir Bychkovsky, Bret Hull, Allen Miu, Hari Balakrishnan, and Samuel Madden. A measurement study of vehicular internet access using in situ Wi-Fi networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, 2006.

[74] Janet Adams. Adaptive Buffer Power Save Mechanism for Mobile Multimedia Streaming. 2007.

[75] Y Agarwal, R Chandra, A Wolman, P Bahl, K Chin, and R Gupta. Wireless wakeups revisited: energy management for voip over wi-fi smartphones. *Proc. of MobiSys*, 2007.

[76] Ashima Gupta and Prasant Mohapatra. Power Consumption and Conservatino in WiFi Based Phones: A Measurement-Based Study. *Proc. of SECON*, 2007.

[77] V Namboodiri and L Gao. Can we reduce Wi-Fi energy consumption during VoIP calls? *Proc. of CoNEXT*, 2007.

[78] E Shih, P Bahl, and MJ Sinclair. Wake on wireless: An event driven energy saving strategy for battery operated devices. *Proc. of MobiCom*, 2002.

[79] P Bahl, A Adya, J Padhye, and A Walman. Reconsidering wireless systems with multiple radios. *Proc. of ACM SigComm*, 2004.

[80] T Pering, Y Agarwal, R Gupta, and R Want. Coolspots: Reducing the power consumption of wireless mobile devices with multiple radio interfaces. *Proc. of MobiSys*, 2006.

[81] IEEE 802.15.4. *Wikipedia*.
`http://en.wikipedia.org/wiki/IEEE_802.15.4`.

[82] First secure Android mobile devices with NFC & Zigbee technologies.
`http://www.taztag.com/index.php`.

[83] CC2530 - A True System-on-Chip Solution for 2.4-GHz IEEE 802.15.4 and ZigBee Applications. *Texas Instruments*.
`http://focus.ti.com/lit/ds/symlink/cc2530.pdf`.

[84] J Liu and L Zhong. Micro power management of active 802.11 interfaces. *Proc. of MobiSys*, 2008.

[85] E Rozner, V Navda, R Ramjee, and S Rayanchu. NAPman: Network-Assisted Power Management for WiFi Devices. *Proc. of MobiSys*, 2010.

[86] Shiao-Li Tsao, Chung-Huei Huang, and Tzu-Ming Lin. Energy conserving packet transmission schemes for video and voice over WLAN. *Proc. of CCNC*, 2006.

[87] Enhua Tan, Lei Guo, Songqing Chen, and Xiaodong Zhang. PSM-throttling: Minimizing Energy Consumption for Bulk Data Communications in WLANs. *Proc. of ICNP*, 2007.

[88] S Singh and CS Raghavendra. PAMAS—power aware multi-access protocol with signalling for ad hoc networks. *Proc. of CCR*, 1998.

[89] P Shenoy and P Radkov. Proxy-assisted power-friendly streaming to mobile devices. *Proc. of MMCN*, 2003.

[90] R BALAKRISHNAN. Minimizing Energy for Wireless Web Access with Bounded Slowdown. *Proc. of MobiCom*, 2002.

[91] C Hu, R Zheng, JC Hou, and L Sha. A microscopic study of power management in IEEE 802.11 wireless networks. *International Journal of Wireless and Mobile Computing*, 2006.

[92] B Gleeson, D Picovici, R Skehill, and J Nelson. Exploring power saving in 802.11 VoIP wireless links*. *Proc. of ICCMC*, 2006.

[93] F LM and M Nilsson. Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking. *Proc. of InfoCom*, 2001.

[94] S Nedevschi, L Popa, G Iannaccone, S Ratnasamy, and D Wetherall. Reducing network energy consumption via sleeping and rate-adaptation. *in Proc. of NSDI*, 2008.

[95] V Namboodiri and L Gao. Towards energy efficient VoIP over wireless LANs. *Proc. of MobiHoc*, 2008.

[96] N Balasubramanian, A Balasubramanian, and A Venkataramani. Energy consumption in mobile phones: a measurement study and implications for network applications. *Proc. of ACM IMC*, 2009.

[97] A Rahmati and L Zhong. Context-for-wireless: context-sensitive energy-efficient wireless data transfer. *Prof. of MobiSys*, 2007.

[98] Y Agarwal, S Hodges, R Chandra, J Scott, P Bahl, and R Gupta. Somniloquy: augmenting network interfaces to reduce PC energy usage. *Proc. of the USENIX NSDI*, 2009.

[99] Ganesh Ananthanarayanan and Ion Stoica. Blue-Fi: enhancing Wi-Fi performance using bluetooth signals. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, 2009.

[100] Y Agarwal, R Gupta, and C Schurgers. Dynamic power management using on demand paging for networked embedded systems. In *Design Automation Conference, 2005. Proceedings of the ASP-DAC 2005. Asia and South Pacific*, 2005.

[101] microSD ZigBee card. *Spectec Computer Co.*
`http://www.spectec.com.tw/sdz537.htm`.

[102] Aditya Akella, Glenn Judd, Srinivasan Seshan, and Peter Steenkiste. Self-management in chaotic wireless deployments. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, 2005.

[103] Allen Miu, Hari Balakrishnan, and Can Emre Koksal. Improving loss resilience with multi-radio diversity in wireless networks. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, 2005.

[104] Ranveer Chandra, Paramvir Bahl, and Pradeep Bahl. MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card. *Proc. of InfoCom*, 2004.

[105] Virtual WiFi. *Microsoft Research*.
`http://research.microsoft.com/en-us/um/redmond/projects/virtualwifi/`.

[106] R. Steward. Stream control transmission protocol. *IETF RFC 4960*, 2007.
`http://tools.ietf.org/html/rfc4960`.

[107] Alan Ford, Costin Raiciu, Mark Handley, and Olivier Bonaventure. Tcp extensions for multipath operation with multiple addresses. *Internet-Draft*, March 2011.
`http://tools.ietf.org/html/draft-ietf-mptcp-multiaddressed-06`.

[108] D Wischik, C Raiciu, and A Greenhalgh. Design, implementation and evaluation of congestion control for multipath TCP. *Proc. of NSDI*, 2011.

[109] Hung-Yun Hsieh and Raghupathy Sivakumar. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, 2002.

[110] Hung-Yun Hsieh, Kyu-Han Kim, Yujie Zhu, and Raghupathy Sivakumar. A receiver-centric transport protocol for mobile hosts with heterogeneous wireless interfaces. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, 2003.

[111] L Magalhaes and R Kravets. Transport level mechanisms for bandwidth aggregation on mobile hosts. In *Network Protocols, 2001. Ninth International Conference on*, 2001.

[112] Božidar Radunović, Christos Gkantsidis, Dinan Gunawardena, and Peter Key. Horizon: balancing tcp over multiple paths in wireless mesh network. In *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*, 2008.

[113] Sébastien Barré, Christoph Paasch, and Olivier Bonaventure. MultiPath TCP: from theory to practice. In *NETWORKING'11: Proceedings of the 10th international IFIP TC 6 conference on Networking*, 2011.

[114] Costin Raiciu, Sébastien Barré, Christopher Pluntke, Adam Greenhalgh, Damon Wischik, and Mark Handley. Improving datacenter performance and robustness with multipath TCP. In *SIGCOMM '11: Proceedings of the ACM SIGCOMM 2011 conference*, 2011.