International Journal of Technology and Engineering System (IJTES)
Vol 7. No.1 2015 Pp. 26-32

available at : www.ijcns.com
ISSN: 0976-1345

A CENTRALIZED APPROACH FOR AUGMENTATION OF ENERGY CONSUMPTION IN RESIDENTIAL BROADBAND GATEWAYS

B.C.PRAVEENA KUMARY[1], Mrs.CHETTIYAR VANI VIVEKANAND[2],
[1,] *PG Student, Cape Institute of Technology,*
email id-chandrikapraveena@gmail.com,
[2,] *Assistant Professor, ECE Department, Cape Institute of Technology*

ABSTRACT

WI-FI interface is known to be a primary energy consumer in mobile devices. The broadband gateway users have long been suffering from high energy consumption due to residential Wi-Fi networks. The residential broadband gateways are significant contributors to overall network energy consumption due to large deployment numbers. The main idea behind this paper is to prototype an architecture which minimizes the power consumption of home gateways. Moreover, home gateways are typically always on, so as to provide continuous online presence to house hold devices such as smart metering, security surveillance, medical monitoring etc.. Hence this architecture advocate a centralized approach, whereby a single authority coordinates the home gateways to maximize energy savings in a fair manner. Reducing the energy consumption of home gateways is to leverage the overlap of Wi-Fi network and aggregate user traffic on fewer gateways, thus by putting remaining to sleep. This centralized control of Wi-Fi network allots bandwidth to the gateways that needs service. This solution can be implemented to heterogeneous internet service providers and also it avoids the client side modifications. That is, there is no need for manual control of broadband gateways. This prototype permits explicit control or session migration. The system is prototype on commodity Wi-Fi access points and tested. By this architecture the energy consumption of broadband gateways can be reduced greatly.

***Index Terms*—Home WiFi networks, energy consumption, bandwidth aggregation, centralized control, system evaluation.**

1.INTRODUCTION

Majority of the power is consumed in the access network [2], due to the sheer volume of user premises equipment (typically a home gateway comprising a modem, router, and wireless access point) and their high per-bit energy consumption. While it is conceivable that future generations of home gateways will implement sleep-on idle (SoI) capability, this will prove ineffective when the household has devices that generate continuous light traffic. A typical household today is estimated to have between 4 and 7 wireless devices, and this number is estimated to grow to 15 within a few years. Devices for VoIP, continuous health-monitoring, smart-metering, security surveillance, etc. are expected to proliferate that require continuous online presence – this "insomnia in the access" [4] has been shown to severely limit the benefits of SoI in home gateways.

This work investigated the feasibility of reducing the combined energy footprint of home gateways by pooling their wireless resources and dynamically aggregating user traffic on to a subset of gateways. The density of WiFi access points in urban areas is known to be high. It is also known that the average load on an access point is quite low. This would suggest that in theory it should be feasible to dynamically sleep/wake access points to save energy while still providing connectivity (and adequate bandwidth) to clients.

There is only one existing proposal for greening residential WiFi networks [4] which takes a distributed approach and embeds intelligence in clients to dynamically aggregate their traffic to a reduced set of APs. While the solution is novel, it is impractical for large-scale adoption as it (a) requires complex client-side machinery (including interface virtualization, reverse NAT, traffic snooping, etc.), which imposes a heavy burden on users and does not extend easily to diverse client platforms, (b) falls apart when the network includes noncompliant (careless or malicious) clients, and (c) does not address fairness (in benefits and costs) which is often topmost on participants' minds.

Unlike the approach in [4], this paper propose a solution architecture in which centralized control is exercised to achieve network energy savings in a fair manner. The centralized entity could be an ISP or any over-the-top third party. Centralization comes with the benefit of realizing optimal or near-optimal solutions more easily, and fairness can be readily incorporated. Centralization does introduce concerns around failures and attacks. However, these issues are better handled via software in a centralized way rather than via embedded distributed protocols. Further, centralization removes a major barrier to adoption, since users need only sign-up and forget, leaving the burden of greening the access network to the centralized entity. There is (inevitably) a price to be paid for client transparency – session migrations from one AP to another cannot be seamless and the transient disruptions during hand-off will cause some traffic loss affecting user quality-of-experience. However, the central controller has visibility into client traffic rates and can therefore choose to perform migrations during periods of "light" traffic, thus minimizing user disruptions. The specific contributions of this paper are:

1) Developed a solution architecture for energy savings in today's residential WiFi networks. This architecture is very suitable for immediate large-scale deployment since it: (a) centralizes control, enabling for optimality and fairness in energy savings, (b) overlays on today's networks at very low cost, allowing for heterogeneity amongst households, (c) has low barrier to entry, as no user side management or client modification is required, and (d) enables substantial energy savings with minimal impact on user performance.
2) Developed (optimal and heuristic) algorithms that maximize energy savings while enforcing fairness, and demonstrate their performance, showing how the energy savings vary with algorithm parameters such as fairness weights and client disruption thresholds.

2 SYSTEM ARCHITECTURE

A central goal of this solution is to minimize the burden on users to participate in the energy savings scheme. To this end, a "setand-forget" approach whereby users opt for the greening service, hand over control of their AP to the green network operator, configure one new SSID on their client devices, and then use their devices as normal. By not requiring users to install and maintain special software on their devices, the scheme is largely transparent to users, encompasses all current and future generation of client WiFi devices, and lowers management, maintenance and support costs for the operator.

The "green operator" may be an ISP or a third party referred to as "operator" takes on the responsibility of greening the aggregate of residential WiFi networks. This operator has management control of all the residential gateways (APs), and makes centralized decisions of their states (awake versus asleep), as well as the corresponding client associations. In order to do this the operator installs the following capabilities at the APs:

Dual-SSIDs: To delineate "home" clients from "guest" clients, each AP is configured with two SSIDs. The "home" SSID is managed by the user (much like today), and can have name and security settings as desired by the user. Devices that connect to the home SSID are identified by the operator as belonging to that user. The "guest" SSID is common across all APs, and is configured with a common security key and traffic on this SSID is tightly controlled by the operator. All user clients have to be configured with both SSIDs.

Client monitoring: The AP has the capability to monitor client activity (traffic flows and their rates) and report these back to the central controller. This information is used by the operator to make the energy savings decisions. The implementation runs an open source IP traffic logging tool called RFlow (an alternative to the industry standard NetFlow application [10]) at the APs to collect this information. Further, the AP also collects information (by snooping on all channels) about clients that are within range but connected to other APs. This information is used by the operator to deduce feasible alternate paths for each client if migration is necessary.

Radio management: The operator turns the radio in the home gateway on/off remotely (over the WAN) to save energy.

Remote management is commonly available on today's gateways. In the future, we hope that gateways will have sleepon-idle (SoI) and remote wake-up capability, which will allow the entire gateway (and not just the radio) to be put to sleep and woken up to enable greater energy savings.

The central controller runs the energy savings algorithm periodically to determine the set of APs that need to be on, and the client associations. The algorithm takes as input the set of client connections and their data rates, the alternate paths available to clients, along with running estimates of costs borne by the APs. It then computes the best set of APs to be kept on, so as to minimize energy, while limiting client traffic disruptions as well as enforcing fairness. The controller then migrates clients between appropriate APs as needed by white/black-listing their MAC addresses at the corresponding APs.

Fairness is needed to ensure that the benefits of energy savings, and the potential data costs of guest client downloads, are shared equally amongst the APs. Centralization makes this relatively easy to achieve. For each AP, the central controller maintains a running cost of supporting “guest” clients. This cost comprises two parts: the *energy cost* is attributable to the power consumption of the AP, and is the sum of a base (static) power that the AP consumes when sleeping (i.e. radio is off), and the dynamic power that is the increment when the AP is active (i.e. radio is on), while the *data cost* is associated with downloads performed by guest clients. Note that home client downloads (i.e. when a client is in its home network) do not count towards this cost, since these are not shared contributions towards the scheme. The guest data cost is balanced across all the users without regard to the volume of each home client download. The combined cost (of energy and data) is averaged for each AP using an exponential averaging technique. This averaged cost is used by the algorithm as a selection weight so that an AP that has incurred high guest cost is more likely to be put to sleep to achieve fair distribution of costs.

Migrations of clients are essential, so as to permit the dynamic aggregation of traffic to a reduced number of gateways for energy savings. Unfortunately, unlike in an enterprise environment, migrations entail disruptions in a residential setting, since the client has to obtain a new IP address associated with the AP (and ISP) it migrates to. It is possible to overcome this problem using sophisticated techniques. The migration is tightly controlled by the central controller by blacklisting the client’s MAC address from the AP it needs to be migrated out of, and white-listing it at the AP it needs to migrate to. This allows the controller to track ISP and power costs incurred by guest clients (on the guest SSID).

3 . OPTIMIZATION AND ALGORITHM

Let us discuss the algorithm for energy savings. First develop an optimization framework, and then develop a heuristic algorithm with known bounds.

3.1 . Optimization Framework

The framework determines the *minimum* set of APs that are needed to provide Internet connectivity to a set of (residential) end-user clients. The APs are numbered 1 to N, and each AP has broadband (i.e. WAN) download bandwidth capacity C_j

Mbps for $j = 1,...,N$. Let $U = \{1, 2, 3,...,n\}$ denote the set of clients, and each client has bandwidth requirement (measured over a certain interval) of b_i Mbps. Denote by e_{ij} the indicator variable that is 1 if and only if client i can connect to AP j on any frequency channel with a specified minimum signal strength (chosen so that the connection is at an acceptable rate, as discussed in §5). Denote by W_j the weight associated with AP j (as briefly described above, these weights are used to control fairness in energy consumption and bandwidth costs across the APs). The optimization framework takes as input the b_i’s, C_j’s, e_{ij}’s and W_j’s as defined above, and computes $\forall i,j$: x_{ij} which is 1 if client i is connected to AP j, and 0 otherwise. Also denote by X_j the binary variable indicating if AP j is on or off. The objective function of this optimization is to minimize the weighted power P:

$$P = \sum_j W_j X_j \quad (1)$$

subject to the following constraints:

$$x_{ij} \leq e_{ij} \quad \forall i,j \quad (2)$$

$$\sum_j x_{ij} = 1 \quad \forall i \quad (3)$$

$$\sum_i b_i x_{ij} \leq C_j, \quad \forall j \quad (4)$$

$$\sum_i x_{ij} \leq X_j \sum_i e_{ij} \quad \forall j \quad (5)$$

Constraint (2) ensures that a client i connects to AP j only if the link is feasible. Constraint (3) restricts each client to connect to exactly one AP, while (4) ensures that the sum of the bandwidth requirements of all clients does not exceed the broadband capacity at the AP. Constraint (5) mandates that an AP is turned on only if at least one client connects to it. Call this problem as minimum weighted set of APs (MWS-AP) that provide network connectivity to all clients.

3.2. NP-Completeness and Heuristic Algorithm

Nnow formally prove that the decision version of the above MWS-AP problem is NP-Complete:

Theorem 1**.** Determining if there exists a solution to the MWSAP problem above, with weighted power no more than P, is NPComplete.

Proof: Given an association of clients to APs (i.e. x_{ij}'s), it is easy to verify that it is feasible (2), connects each client to exactly one AP (3), and does not exceed the bandwidth capacity of an AP (4). APs that connect to at least one client need to be on (5), and their weighted power is verified to be at most P. Therefore MWS-AP is in NP.

To prove that MWS-AP is NP-Hard, show that the Weighted Set Cover problem can be reduced to MWS-AP in polynomial time. Let the universe of elements be $U = \{1, 2, 3, n\}$, where each element corresponds to a client. Let $S = \{S_1, S2, SN\}$ denote the collection of subsets of U, where $S_j \subset U$ corresponds to the clients that can connect to AP j. Let W_j be the weight of set S_j, corresponding to the weight of AP j. The above (trivial) transformation is polynomial in time, and it is easy to see that if the MWS-AP (with capacity C_j of each AP set to ∞) can be solved in polynomial time, this will directly yield a solution to the weighted set cover problem.

The best known algorithm for weighted set cover has an approximation factor of $O(\ln n)$, where n is the number of elements in the universe. Adapt this greedy algorithm for MWS-AP by selecting at each step the AP that has the highest value of the number of (as yet unconnected) clients divided by the weight of the AP. The algorithm is formally described next.

Algorithm 1 Determine the set of active APs

Inputs: Set of clients U, visible client set Sj and weight Wj for each AP j

Output: Set I of APs that should be active Temporary variables: X, I

1: X ← U, I ← ϕ

2: Repeat until X = ϕ

Pick AP j with smallest |

3: $I \leftarrow I \cup \{j\}$, $X \leftarrow X \setminus S$ X ∩ Sj| /Wj

4: j

5: Output I

The algorithm above takes as input the set of clients U, and for each AP j, its coverage of clients Sj and weight Wj. Internal variable X keeps track of the set of clients that are uncovered, and is initialized to the entire set of clients U in step 1. Variable I stores the selected APs, and is initialized to the null set in step

1. The algorithm operates in a loop till all clients are covered, i.e. X = ϕ (step 2). In each iteration, the AP j which has the maximum ratio of unconnected clients (|X ∩ Sj|) to weight (Wj) is selected (step 3). This AP is added to the set I of selected APs and the clients Sj it covers are removed from X (step 4). The set of active APs I is output in step 5.

3.3. Client Migrations and AP Fairness

This algorithm determines the set of APs that need to be on, but does not clarify the attachment of clients to the APs. Therefore refine the algorithm in the following ways:

For each "active" client, i.e. one whose traffic intensity is above a certain threshold (we will study the impact of this threshold in the next section), do not want to disrupt its existing connectivity as it can impact user quality-of-experience. The algorithm therefore begins by keeping such APs on, i.e. in step 1 it will initialize the set I to include such APs, and the set X to exclude clients covered by these APs.

The algorithm outputs the set of APs that need to be on. A client may have multiple APs within range that are on, and use the following method to decide which one it connects to: (a) if the client's home AP comes on, it is migrated to its home SSID (as it is unnecessary to incur a guest cost when the home is available), (b) otherwise, if the client's current AP is on, the client stays connected to the same AP and is not migrated (so as to minimize migration disruptions), and (c) if neither the home AP nor the current AP of the client are on, the client is migrated to the AP with least weight (for fairness, as described next).

4 . RELATED WORK

Nnow provide a brief overview of recent work in the context of reducing energy consumption of access networks, centralized control and AP-level coordination of WiFi networks.

Wire line access networks: Bianco et al. [21] present trends from a Telco-perspective (Telecom Italia) and discuss the importance of minimizing energy consumption for enabling sustainable next generation FTTx access networks. Greening

DSL networks has also received widespread attention – Tsiaflakis et al. [22] and Guenach et al. [23] revisit the dynamic spectrum management problem in DSL systems and make it power-aware by incorporating constraints for limiting the transmit power. Their results indicate that there exists a tradeoff between power savings and data rates. The problem of (re)designing DSL networks to minimize power consumption is studied by Bhaumik et al. [24] who show that the energy efficiency of access networks can be improved by replacing large monolithic DSLAMs with small DSLAM units closer to the customer.

Mobile devices: Based on the observation that WiFi radios in smartphones consume significant amount of power when active, Rozner et al. [25] propose NAP man, a system to minimize WiFi energy consumption in mobile devices. Although PSM (Power Save Mode) is part of the WiFi standard, the authors note that competing background traffic can adversely impact energy consumption. Consequently, an energy-aware fair scheduling algorithm is developed yielding up to 70% savings. Manweiler et al. [26] argue that NAP man is most effective when a single access point is connected to multiple clients. Since in practice a client is within the range of multiple access points, this scenario strongly impacts energy consumption. A new system called Sleep well is designed and savings of up to 51% is demonstrated. Finally, Zhang et al. [27] found that over 60% of the WiFi energy is consumed when the device is in the idle listening (IL) state, even in the presence of PSM. While prior work minimizes the duration a client spends in IL, their proposal E-MiLi reduces the power consumption of IL by reducing the clock rate of the radio when the device is in IL. Results show savings of ≈ 44%.

Wireless access networks: Jardosh et al. [8] and Goma [4] et al. demonstrate greening enterprise and residential WiFi networks respectively. However as argued in §1, the former does not apply to the residential setting because it does not address heterogeneity and fairness issues, while the latter uses a distributed approach and places the onus on end-users, which is fundamentally different from proposed technique that is centralized and controlled by an operator.

Aggregating bandwidth: As the energy saving schemes primarily rely on bandwidth from neighbouring wireless gateways being pooled and shared across many users, now briefly mention few works that achieve this aggregation. Jakubczak et al. [7] and Tan et al. [28] propose Link-alike and CUBS respectively to effectively aggregate and share the uplink bandwidth of neighbouring gateways. Their results indicate that the throughput of upload intensive applications improve by more than 30%. Giustiniano et al. [29] show that fairness is an important metric that needs to be factored in when aggregating capacity from multiple access points, which was not considered by Kandula et al. [30]. As a solution, they propose a system called THEMIS that incorporates fairness and addresses the problem in a distributed manner without requiring any changes to the network. COMBINE by Ananthanarayanan et al. [31] is a system that significantly improves the download speed at any node by pooling together bandwidth from the wide area network. Open Radio [32], based on the software defined networking paradigm, advocates pooling together resources from various wireless access technologies, i.e. WiFi, 3G, LTE etc., for improved network performance and ease of management. It is important to note that while the above works aim to maximize system throughput, reducing the energy consumption has not been their primary focus.

Centralized control and AP coordination: Centralized control has been efficiently leveraged in enterprise WLANs for fault management [33], power control and channel assignment of APs [34], and client localization [35]. Srivastava et al. [36] develop a framework called CENTAUR for mitigating hidden or exposed terminal

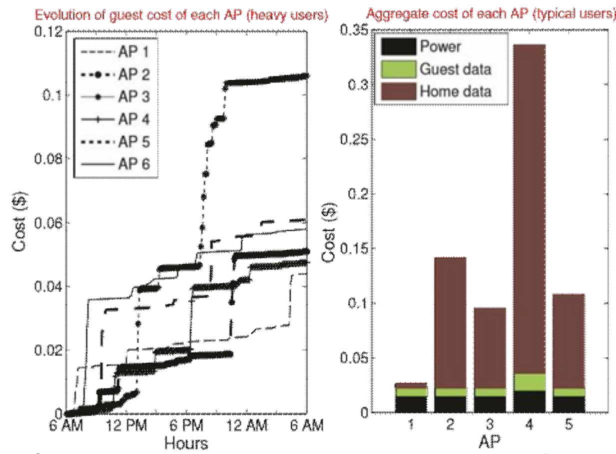

interference in enterprise WLANs. CENTAUR is shown to improve the performance of the network substantially due to improved utilization of the wireless medium. Much like our scheme, the centralization functionality of CENTAUR is implemented in a single central server and requires minor modifications to the APs. No changes are needed to the end-user client devices. A system called FLUID was built by Rayanchu et al. [37] that also employs a central controller and assigns the centre frequencies and channel widths to the APs dynamically, based on traffic demand. Flexible channelization combined with data scheduling is shown to further improve performance. Experimental results demonstrate that FLUID improves the median throughput by 59 %. Akella et al. [38] develop an algorithm called Power controlled Estimated Rate Feedback (PERF) for improving end-user performance in chaotic (i.e. unplanned and unmanaged) dense WiFi network deployments. Home WiFi networks belong to the category of chaotic networks. The key idea behind PERF lies in automatic management of transmission power levels and rates of APs and clients so that interference between neighbouring APs is minimized. Experimental results highlight the benefits of PERF; a client of a highly utilized AP located next to another such pair sees a 20-fold increase in throughput. In essence, this work shows that it is possible to improve the performance of WiFi networks if the APs can be made to coordinate with each other. In the context of home networks, Patro et al. [39] deploy an infrastructure called WiSe to measure and monitor the performance of home WiFi networks. The APs are configured with specialized software and communicate with WiSe's measurement controller using open APIs. Analysis of traces from 30 households spanning over 6 months show that a majority of the links performed well but poor quality was observed about 2% of the time. Finally, Manweiler et al. [40] develop a tool called RxIP for monitoring the health of residential WiFi networks. Specifically, the APs announce their IP addresses periodically. Neighbouring APs that receive these messages then relay them so that they are received by potential hidden terminals. As new APs are discovered, neighbouring APs establish a control channel with these APs over the wired Internet, thus enabling coordination between APs. Results show an improvement of 57% in the median throughput in symmetric hidden terminals.

7 . CONCLUSIONS

Residential broadband access gateways are a major contributor to overall network energy consumption due to their widespread deployment. This paper proposed, evaluated and prototyped a scheme for aggregating users on to a fewer set of WiFi access points to reduce energy consumption. Also made the following contributions: (1) Developed a centralized architecture that works across heterogeneous ISPs and clients, and allows for fairness in energy savings, (2) Developed (optimal and heuristic) algorithms, and also studied the trade-off between energy savings and session disruptions using campus WiFi traces. This scheme is centralized putting the burden on the operator rather than the user, thereby significantly reducing the barrier-to adoption for wide-scale deployment.

REFERENCES

R. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. Sorin. Energy Consumption of IP Networks. In Proc. European Conf. on Optical Communications, Belgium, Sep 2008.

M. Pickavet. Energy in ICT – Trends and Research Directions. In IEEE Advanced Networking and Telecommunications Systems (ANTS), India, Dec 2009.

Number of housing units in California in 2011. http://quickfacts. census.gov/qfd/states/06000.html.

E. Goma et al. Insomnia in the Access or How to Curb Access Network Related Energy Consumption. In Proc. ACM SIGCOMM, Canada, Aug 2011.

BeWiFi from Telefonica. http://www.bewifi.es/.

Xfinity from Comcast. http://www.comcast.com/.

S. Jakubczak et al. Link-alike: Using Wireless to Share Network Resources in a Neighbourhood. ACM SIGMOBILE MC2R, 12(4):798–814, Oct 2008.

A.P. Jardosh et al. Green WLANs: On-Demand WLAN Infrastructures. Springer Mobile Networks and Applications (MONET), 14(6):798–814, Dec 2009.

DD-WRT. Open source Linux-based firmware. www.dd-wrt.com.

Cisco IOS NetFlow. www.cisco.com/go/netflow.

Wiviz. www.dd-wrt.com/wiki/index.php/Wiviz.

TP-LINK WAP. www.tp-link.com/en/products/details/?model= TLWR1043ND.

R. Murty, J. Padhye, R Chandra, A. Wolman, and B. Zill. Designing High Performance Enterprise Wi-Fi Networks. In Proc. USENIX NSDI, USA, Apr 2008.

M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In Proc. IEEE INFOCOM, USA, Mar/Apr 2003.

Sandvine. Global Internet Phenomena Spotlight: Netflix Rising. Technical report, May 2011.

T. Beauvisage. Computer Usage in Daily Life. In Proc. ACM Computer Human Interaction (CHI), USA, Apr 2009.

N. Burton, M. Haynes, J. Uffelen, W. Brown, and G. Turrell. MidAged Adults Sitting Time in Three Contexts. American Journal of Preventive Medicine, 42(4):363–373, 2012. [18] Broadband data block cost from Internode. http://www. internode.on.net/residential/product features/data blocks.

[19] D. Rossi, M. Mellia, and M. Meo. A Detailed Measurement of
Skype Network Traffic. In Proc. USENIX IPTPS, Tampa, Florida,
USA, Feb 2008.

[20] Flow-tools for NetFlow. www.splintered.net/sw/flow-tools/docs/
flowtools.html.

C. Bianco, F. Cucchietti, and G. Griffa. Energy Consumption Trends in the Next Generation Access Network: A Telco Perspective. In Proc.
INTELEC, Italy, Sep/Oct 2007.

P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen. Green DSL: EnergyEfficient DSM. In Proc. IEEE ICC, Germany, Jun 2009.

M. Guenach, C. Nuzman, J. Maes, and M. Peeters. On Power Optimization in DSL Systems. In Proc. GreenComm, Germany, Jun 2009.

S. Bhaumik, D. Chuck G. Narlikar, and G. Wilfong. Energy Efficient Design and Optimization of Wire line Access Networks. In Proc. IEEE INFOCOM Mini-Conference, China, Apr 2011.

E. Rozner, V. Navda, R. Ramjee, and S. Rayanchu. NAPman: NetworkAssisted Power Management for WiFi Devices. In Proc. ACM MobiSys, USA, Jun 2010.

J. Manweiler and R. R. Choudhury. Avoiding the Rush Hours: WiFi Energy Management via Traffic Isolation. In Proc. ACM MobiSys, USA, Jun/Jul 2011.

X. Zhang and K. G. Shen. E-MiLi: Energy-Minimizing Idle Listening in Wireless Networks. In Proc. ACM MOBICOM, USA, Sep 2011.

E. Tan, L. Guo, S. Chen, and X. Zhang. CUBS: Coordinated Upload Bandwidth Sharing in Residential Networks. In Proc. IEEE ICNP, USA, Oct 2009.

D. Giustiniano, E. Goma, A. Lopez Toledo, I. Dangerfield, J. Morillo, and P. Rodriguez. Fair WLAN Backhaul Aggregation. In Proc. ACM MOBICOM, USA, Sep 2010.

S. Kandula, K. C.-J. Lin, T. Badirkhanli, and D. Katabi. FatVAP: Aggregating AP Backhaul Capacity to

Maximize Throughput. In Proc. USENIX NSDI, USA, Apr 2008.
G. Ananthanarayanan, V. N. Padmanabhan, L. Ravindranath, and C. A. Thekkath. COMBINE: Leveraging the Power of Wireless Peers through Collaborative Downloading. In Proc. ACM MobiSys, USA, Jun 2007.
Open Radio. http://snsg.stanford.edu/projects/openradio.
Y. Cheng et al. Automating Cross Layer Diagnosis of Enterprise Wireless Networks. In Proc. ACM SIGCOMM, Japan, Aug 2007.