

Towards an Energy-Star WLAN Infrastructure

Amit P. Jardosh[§], Gianluca Iannaccone[†], Konstantina Papagiannaki[†], Bapi Vinnakota[‡]

[§]Dept. of Computer Science, UCSB [†]Intel Research [‡]Intel Corporation
amitj@cs.ucsb.edu, {gianluca.iannaccone, dina.papagiannaki, bapi.vinnakota}@intel.com

Abstract

Mobile users' demand for omnipresent, reliable, and high throughput services from a wireless LAN infrastructure has led a significant number of enterprises and public bodies to deploy large-scale high-density IEEE 802.11-based WLANs. These WLANs feature several hundreds to thousands of APs that are managed centrally and are typically dimensioned for peak user demands. Unfortunately, peak demands are occasional and WLANs often under-utilized. Using measurements from two different WLANs we show that the under-utilization of WLANs leads to a significant energy wastage, primarily due to their always-on nature. We discuss designs that can render such WLANs energy efficient and describe energy star infrastructures that can maximize energy savings in different scenarios. We also propose a simple algorithm, Green-clustering for energy star WLAN infrastructures that can lead to up to 90% energy savings. We argue that current and future large-scale WLANs must integrate energy efficiency as a design constraint. We also discuss avenues of further research on energy-star WLAN infrastructures.

1 Introduction

Large, centrally controlled wireless local area networks (WLANs) consisting of hundreds of access points (APs), switches, and controllers [1, 2, 4, 5, 6] are now commonplace installations at university campuses [14] and corporate buildings [8]. In such centrally controlled WLANs the *thin* APs are responsible for minimal tasks such as beaconing and probe responses, while the majority of the tasks such as association and authentication are performed at the centralized controller. The APs and the controller communicate using GRE or LWAPP tunnels formed over a wired back-haul network consisting mainly of switches. Using these tunnels each AP encrypts and sends all MAC layer frames received from clients to the central controller and the central controller sends MAC-layer frames to each client via the AP it is associated with [3].

Compared to distributed WLAN APs, APs in centrally controlled WLANs are much simpler and therefore, cheaper. Also, since their centralized controllers maintain global knowledge about the network state, managing and securing WLANs is a relatively easier task. This ease of deployment and the increasing user demands for high-capacity WLANs has catalyzed the emergence of several companies [1, 2, 4, 5, 6] that deploy highly dense WLANs. These WLANs feature hundreds or even thousands of APs that provide redundant overlapping coverage to meet the capacity demands and allow fault-tolerance. For instance, a particular corporate WLAN features more than 5000 APs [8], while the Dartmouth College campus WLAN has 1300 APs since 2005.

The obvious advantage of such large-scale redundant WLANs is that the peak capacity they are designed for is sufficient to provide users with bandwidth for most interactive and throughput-intensive applications such as multimedia voice, video streaming, and online games. The downside is that peak usage times are rare and sometimes even isolated to a portion of the WLAN. As a result, the energy cost of maintaining hundreds to thousands of always-on idle APs and wired back-haul switches is significant. This observation raises a serious concern about the energy wastage in hundreds of such large-scale WLANs across the world. As WLAN companies procure new enterprise customers each month [1, 2, 4, 5, 6] and the average size of WLANs almost doubles each year [14, 9], the energy problem is bound to escalate.

Based on these concerns, we strongly advocate the use of energy-efficiency as a critical design constraint for current-day and future large-scale WLANs. In this paper, we propose that WLANs be re-designed so that the power consumed by the WLAN *scales* with its offered load; however, without compromising WLAN design requirements regarding coverage, client QoS, and redundancy.

To validate our claim that large-scale WLANs are indeed under-utilized and peak usage times are occasional, we utilize real WLAN data traces from two different WLANs: a 125-AP single building corporate WLAN at Intel Corporation, Oregon, and a 500-AP campus-wide WLAN at Dart-




Figure 1. A centralized WLAN infrastructure.

mouth College [15]. We then propose an elegant and effective algorithm, called *Green-clustering* which strategically powers on and off WLAN equipment (i.e., APs, switches, and controllers) based on user demands. We use green-clustering as a tool to demonstrate that each WLAN is a candidate for significant energy savings, although the solution may differ based on their geographical spread, logistical constraints, and design choices.

To our knowledge, this is the first paper that brings forth awareness about the potential energy wastage in large-scale WLANs and also discusses the solution space with a list of requirements and design choices. Energy wastage has been previously cited as a serious concern in large data-centers [11] and in the wired Internet [12]. Both these scenarios target solutions that power on and off equipment based on local observation of either processing load or the arrival of data packets, respectively. However, the powering on and off of APs and switches in a WLAN is a harder problem to solve because of the spatially dispersed nature of WLANs and the varying wireless channel propagation characteristics. We believe that such design challenges and the diverse deployment characteristics of WLANs makes the design of effective energy-efficient solutions a challenging task. Our solution is the first step towards this effort.

The rest of the paper is organized as follows. In Section 2 we discuss the energy requirements in high-density WLANs. In Section 3 we present two WLAN usage studies and evaluate energy wastage in each. Section 4 lists the requirements from an energy-star WLAN infrastructure and the design choices that need to be made. In Section 5 we propose and evaluate a simple algorithm called green-clustering and discuss further extensions to it, and in Section 6 we present our conclusions.

2 Energy Consumption in High Density WLANs

Figure 1 shows a sample architecture of a centralized large-scale WLAN. The APs and the controller are connected via a single or multiple layers of wired switches.

The number of switches and the connections between the switches are determined based on the geography of the enterprise and entailing logistical constraints. These switches also typically power the APs via IEEE 802.3af *Power-over-Ethernet* (PoE) ports.

In a centrally controlled WLAN, the two main consumers of energy are: thin APs and the switches they are connected to. Each thin AP draws power from PoE ports on PoE-compatible switches. Although the PoE specifications suggest a maximum power consumption of 15.4W, each AP typically draws up to 10 W [1, 2, 4, 5, 6]. Based on the US electricity costs of about 15 cents per killo-watt-hour (kWh) [7], every 100 APs in a WLAN consume up to \$108 worth of power per month. In addition to thin APs, the switches used in centralized WLANs that support 24 to up to 72 APs consume up to 350W per hour each in excess to the power they supply to the APs connected to them [1, 2, 4, 5, 6]. These numbers indicate that the operational cost of current WLANs is far from negligible, a cost that becomes more noticeable once one starts looking at the hundreds to thousands of WLANs deployed worldwide. This wastage is bound to escalate in the future.

3 WLAN Usage Studies

In this section we utilize real data traces from two WLANs - the Dartmouth college WLAN and the Intel corporate WLAN, to study the usage of APs in such scenarios and the potential for energy savings. We use these WLANs because they represent two different design choices. The first WLAN corresponds to the network deployed across one 4-floor building of Intel's campus in Portland, Oregon, and features 125 APs and 1 controller. Each floor features one AP installed for every 4 to 6 cubicles (typical inter-AP distance is approximately 5 meters). Each floor is 80 meters \times 28 meters with a density of 1 AP per 52 m^2 . The main design goal of the Intel WLAN is to provide sufficient capacity and overlapping coverage to corporate users so that they can accomplish bandwidth-intensive tasks even while on the move. On the other hand, the main design goal of the 500-AP Dartmouth WLAN is to only provide basic connectivity to campus users. This objective is reflected in the smaller AP density of 500 APs in 188 buildings on a 4 Km \times 5 Km campus. Notice that these statements reflect the state of the Dartmouth WLAN in 2002, the year when the publicly available data was collected.

For our study, we use Simple Network Management (SNMP) logs collected at 5-minute intervals from each AP. These SNMP logs provide information on all the users connected to each AP at the time each log is collected. We choose to use a representative single week's data out of the 2 years of publicly available Dartmouth WLAN logs [15] and 2 months of Intel WLAN logs. Using these one-week

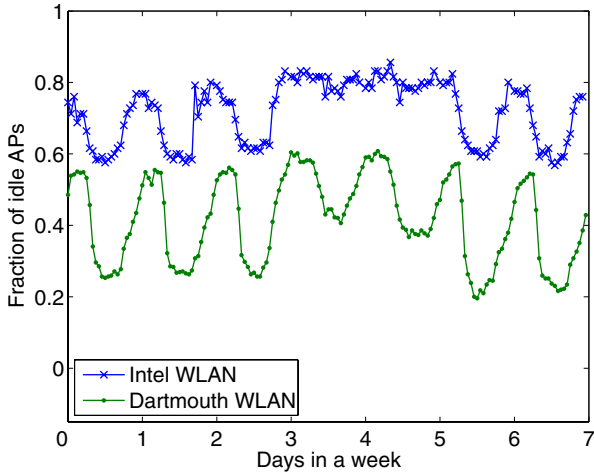


Figure 2. Fraction of idle APs in the Dartmouth and Intel WLANs. Day 4 is a Saturday.

SNMP logs, we identify the APs that are *idle* during an entire one-hour interval; idle APs are APs that do not have even a single user associated with them [10]. Figure 2 shows the fraction of idle APs per hour during the entire week for both WLANs. The figure clearly shows that a large fraction of the APs in both the WLANs remain idle during any one-hour interval. We also observe that the fraction is much larger in the Intel WLAN, and more so during nights and weekends. That is because the Dartmouth WLAN APs are sparsely dispersed and service a greater number of users.

We believe that these results are representative of several other WLAN deployments as well. These results show that WLAN capacity is frequently under-utilized - even in vastly different WLAN scenarios. Idle APs in WLANs across the globe directly equate to enormous wastage of energy that is used to keep the idle APs powered on. Based on these observations, we conclude that *WLANs must be (re-) designed with energy-efficiency as a critical design constraint*. Such energy-efficient or *energy-star* WLAN infrastructures should appropriately power on and off WLAN equipment so that they stand to save both critical energy, as well as monetary resources.

4 Energy-star WLAN Design Requirements

Powering off portions of a WLAN infrastructure has significant benefits. It saves energy, monetary costs, and also reduces the complexity of configuring *all* APs to optimal frequencies and transmit powers [16]. However, the design of such energy-star WLANs should not impact the performance of the WLAN and its users. To avoid such an impact, we present a list of design requirements and strategies on

how they can be adhered to, even while maximizing energy savings.

Coverage: WLAN coverage is defined as the spatial region on which clients can associate with a WLAN. Powering off APs should maintain the original WLAN coverage. To meet this requirement, energy-star WLANs should ensure that there is at least one layer of APs that retains the coverage to the region covered by the always-on WLAN. It is the redundant APs that can be powered off appropriately for energy savings.

Redundancy: WLAN redundancy is defined as the excess number of APs deployed by WLAN administrators to sustain heavy load or to compensate for equipment failure. We argue that in energy-star WLANs, it is necessary and sufficient to maintain the same redundancy *factor* instead of total redundancy. In other words, if an always-on WLAN uses N APs for complete WLAN coverage and for sustaining peak bandwidth demands, and β (≥ 1) is the redundancy factor chosen for the WLAN, the total number of APs deployed will be $\beta \times N$. During off-peak demands, if the WLAN controller determines that only M ($M \leq N$) APs are needed to maintain coverage and support demand, $\beta \times M$ APs can be powered on, such that the same redundancy factor β is maintained. The redundancy factor can also be different for different parts of the same WLAN depending on its original design.

Client QoS: Client QoS is loosely defined as the performance each client can gain from a WLAN, for instance, throughput and end-to-end delay. The goal of an energy-star WLAN architecture is to save energy without jeopardizing client QoS. Client QoS, like the redundancy factor, can be maintained by appropriately powering on a sufficient and necessary number of APs close to the users in the WLAN to ensure that clients experience throughput and end-to-end delays similar to their always-on WLAN counterparts. Since central controllers in energy-star WLANs maintain complete information about client activity, the controller can power on more APs if it detects any degradation in client performance indicated by a drop in data rates or signal-to-noise ratios.

Responsiveness: Immediate response of the WLAN to client requests should be a common feature of both energy-efficient and traditional deployments. Typical boot-up times for APs and switches lie in the order of milliseconds with the potential to reach a few seconds. Consequently, an energy efficient design could potentially lead to delayed responses to user requests for association. However, in energy-efficient WLANs, already on APs could temporarily accommodate increased demand until new APs are powered on. As a result an AP may be required to support a few additional users for a short duration of a few seconds until a new AP is powered on. Moreover, one could also envision solutions where APs and switches are powered on

in advance based on predicted increase in load and user association events.

5 Green-clustering algorithm

In this section we devise a simple algorithm, called *green-clustering*, that adheres to the design requirements discussed in Section 4. The objective of this algorithm is to simply explore the magnitude of energy savings that can be achieved in different WLANs and thus the feasibility of adopting research on energy-star WLANs; we find that our results are highly supportive of this objective. Further exploration and exhaustive evaluation of a mature algorithm forms our future work.

Opportunities for energy savings: Previously proposed distributed approaches for energy efficient WLANs rely on locally derived information. For instance, a distributed approach called Wake-on-WLAN powers on or off APs if a sensor attached to the AP detects the presence or absence of users in the vicinity of an AP [17]. As a result, more than a necessary number of APs may power themselves on when only a single AP may be sufficient.

The green-clustering algorithm enables the central controller in a WLAN to make *smart* decisions to power on and off portions of a WLAN. These smart decisions ensure that all the requirements of an energy-star WLAN design are met while maximizing energy savings. The green-clustering algorithm is a policy-based algorithm. It chooses to power WLAN APs differently based on the deployment and logistical constraints of each WLAN. In the following part of this paper we use the Intel and Dartmouth WLANs as examples of different deployments and show how differently green-clustering can save energy in both these scenarios.

Ensuring WLAN coverage: In green-clustering we form clusters of APs based on the distance between them. The basic premise is that if APs in a cluster are *close enough*, a single AP from each cluster is sufficient to provide basic coverage to users in the vicinity of any AP in that cluster. Such clustering is possible in high-density WLANs because APs are deployed close to each other to provide overlapping coverage and high bandwidth to users in their proximity. Most IEEE 802.11b APs can sustain the highest transmission rate of 11 Mbps for users at a distance of up to 50 meters [13]. If we assume that each AP in a high-density WLAN was originally configured to support a region of radius 20 meters, we can form clusters of APs that are up to 30 meters from each other and such that *any* single AP from that cluster can still support the highest transmission rate for users even at the edge of the cluster region. This single AP is termed as the *cluster-lead AP*. In Figure 3(a), if APs *A* to *E* all lie within a distance of 30 meters from each other, they can all form a single cluster. If AP *E* is chosen as

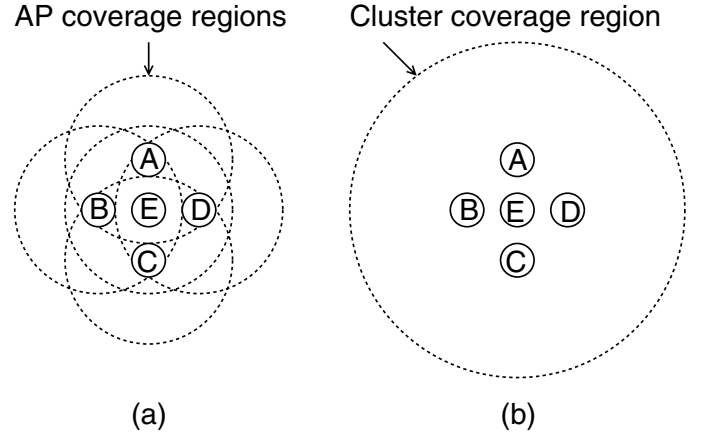


Figure 3. Illustration of green-cluster formation of APs.

the cluster-lead AP, it can provide coverage to a composite region of coverage as shown in Figure 3(b). Using this green-clustering algorithm and the location information of APs on the Intel and the Dartmouth WLANs, we form 9 and 340 clusters of APs and select 9 and 340 cluster lead APs, respectively. Since the Intel WLAN is a more dense deployment of APs than the Dartmouth WLAN, the Intel WLAN consists of fewer clusters but more APs per cluster than the Dartmouth WLAN.

In green-clustering we use euclidian distances between APs to form clusters. However, if radio propagation characteristics of the WLAN coverage area were known via site surveys or computed using innovative measurement techniques, better clustering of APs can be pursued. Better clustering results in better QoS for clients and greater energy savings.

Scaling energy consumption with offered load: The WLAN controller monitors the number of users requesting WLAN access in the vicinity of each cluster. Based on this information, the controller decides to power on and off the other APs or *surrogate APs* within each cluster. Such decisions could also utilize metrics such as aggregate traffic levels, communication data rates, and number of data frame retransmissions. Since the Intel WLAN deploys 1 AP for every 4 users, the green-clustering algorithm uses 4 as the number of simultaneous users per AP to power on an extra surrogate AP. Also, at regular intervals, surrogate APs are powered off if there are less than 4 users per AP within a cluster. Note that such a mechanism may necessitate the re-configuration of AP frequency and transmit power to ensure optimal performance with minimal interference. Such re-configurations can occur on-the-fly using known techniques [16] or different configurations settings can be pre-computed.

The green-clustering algorithm currently uses a fixed number of users as a trigger to power surrogate APs. However, for greater energy savings we believe that accurate user load estimation techniques, efficient user-to-AP assignment strategies, and corresponding optimization algorithms can be developed as future work.

Energy savings via switch reconfigurations: In addition to controlling APs, centralized controllers can choose to power on and off power-hungry WLAN switches. This is possible only if WLAN administrators (re-)configure WLANs such that APs are connected to dedicated switches and AP-switch groups can be powered off simultaneously. These AP-switch groups can be created based on either the functional role of the APs (the level of redundancy provided by each AP) or their temporal usage patterns (APs that experience strong correlation in their usage patterns). Such grouping of APs and switches is possible in many enterprise WLAN installations such as the Intel WLAN where the APs are connected to switches within the same server-room. However, like the Dartmouth WLAN, not all WLAN deployments can allow such groupings if their APs are geographical dispersed and not many APs are connected to the switches to begin with.

For our evaluation, we rank and divide the 125 Intel WLAN APs into 4 sets based on their usage during an entire week. Lower rank APs are APs that are used often or at all times, while the higher rank APs are APs that are rarely on or never used. The 4 32-port switches used in the Intel WLAN are reconfigured so that switches with lower rank APs are likely to remain on often or at all times, and switches with higher ranked APs will occasionally be used, thus adding to energy savings.

Simulation: We use a custom simulator to replay the association of users with APs as extracted from Intel and Dartmouth SNMP logs. If a user is associated with AP i in the original WLAN, then the simulator associates the user with any AP in i 's cluster. If the number of users supported by the cluster exceeds the limit of 4 per AP, then the simulator powers on an extra surrogate AP in the same cluster. In case of the Intel WLAN, the simulator prioritizes the selection of a surrogate AP that is on an already powered on switch over a powered off one, so that more switches can remain powered off for as long as possible.

Savings: We compute per-hour percentage energy savings from the green-clustering algorithm based on the number of APs that remain powered off during the hour. In other words, if n APs out of the total N ($N = 125$ for the Intel WLAN and 500 for the Dartmouth WLAN) remain powered off during an hour, then the percentage power saved over the always-on original WLAN is $(\frac{n}{N} \times 100)\%$. These percentage values are computed for each hour during a period of one week for both the Intel and Dartmouth WLANs. Also, to compute the additional energy savings in the Intel

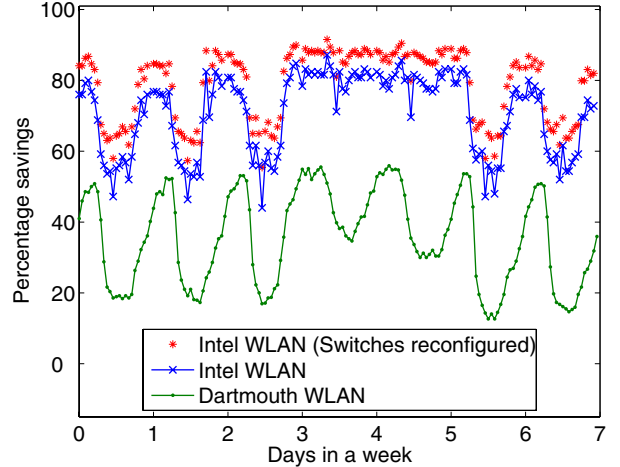


Figure 4. Percentage energy savings in the Dartmouth and Intel WLANs.

WLAN due to switch reconfigurations, we record the number of switches used by the simulator at each hour. If a 32-port switch consumes 320W of power, it is equivalent to the power consumed by 32 typical 10W APs. Therefore, if s switches out of S are powered during an hour, the percentage energy savings are $(\frac{n+(s \times 32)}{N+(S \times 32)} \times 100)\%$.

Figure 4 shows the percentage energy savings from using green-clustering on the Dartmouth and Intel WLANs. The savings vary from 20% to 50% in the Dartmouth WLAN and 50% to 80% in the Intel WLAN. Moreover, during a few 1-hour intervals, up to 90% energy savings are achieved in the Intel WLAN if the switches are reconfigured. The difference in energy savings between the two WLANs occurs because of the much greater number of always-on cluster-lead APs and the significant AP usage resulting from the sparse density of the Dartmouth WLAN. These results show that power *can* be saved in geographically spread Dartmouth-like WLANs also, without WLAN reconfigurations.

6 Conclusion

Using real data traces, this paper shows that large-scale WLANs are indeed under-utilized and result in power wastage in large-scale WLANs of today, which will certainly aggravate in the future. We show that simple algorithms can be used in such WLANs to conserve power, sometimes close to 90%. However, we argue that the algorithms and results presented in this paper are a naive first look at energy efficient wireless WLAN designs. The primary objective of this paper is to promote further research on energy-efficient infrastructure designs. We discuss some

of these extensions in the paper, but foresee many other avenues for further development.

We conclude that *each* large-scale centrally controlled WLAN is a candidate for energy savings, even while meeting its original design requirements. However, the volume of energy savings and the approaches to save power can significantly differ based on the WLAN's deployment characteristics, the number of clients, and their usage characteristics.

WLAN networks are becoming larger and with a higher density of radio devices per user. This trend is gaining momentum in the industry with a large number of vendors providing commercial solutions without any effort towards energy efficiency. We believe that this paper makes a timely proposal for energy-conscious design guidelines in high-density WLANs.

Acknowledgments

We thank Omer Ben-Shalom of Intel Corporation for sharing WLAN data-traces and the anonymous reviewers of this paper for their valuable comments.

References

- [1] Aruba Wireless Networks. <http://www.arubanetworks.com>.
- [2] Cisco Unified Wireless Network. <http://www.cisco.com/go/unifiedwireless>.
- [3] IETF CAPWAP Working Group. <http://www.ietf.org/html.charters/capwap-charter.html>.
- [4] Meru Networks. <http://www.merunetworks.com/>.
- [5] Symbol Technologies. <http://www.symbol.com>.
- [6] Trapeze Networks. <http://www.trapezenetworks.com>.
- [7] U.S. Energy Information Administration. <http://www.eia.doe.gov/cneaf/electricity/>.
- [8] Aruba Selected by Microsoft For Next Generation Wireless LAN. <http://www.arubanetworks.com/news/release/2005/06/13>, 2005.
- [9] Forrester Research. <http://www.forrester.com>, 2006.
- [10] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-area Network. In *ACM Mobisys*, San Francisco, CA, May 2003.
- [11] J. S. Chase, D. C. Anderson, P. N. Thakar, A. Vahdat, and R. P. Doyle. Managing Energy and Server Resources in Hosting Centres. In *SOSP*, Banff, Canada, October 2001.
- [12] M. Gupta and S. Singh. Greening of the Internet. In *ACM SIGCOMM*, Karlsruhe, Germany, August 2003.
- [13] T. Henderson, D. Kotz, and I. Ayzov. The Changing Use of a Mature Campus-wide Wireless Network. In *ACM MOBICOM*, Philadelphia, PA, September 2004.
- [14] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *ACM MOBICOM*, Atlanta, GA, September 2002.
- [15] D. Kotz, T. Henderson, and I. Ayzov. CRAWDAD data set dartmouth/campus (v. 2004-11-09). Downloaded from <http://crawdad.cs.dartmouth.edu/dartmouth/campus/snmp>, 2004.
- [16] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A Client-driven Approach for Channel Management in Wireless LANs. In *Proceedings of IEEE Infocom*, Barcelona, Spain, April 2006.
- [17] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-WLAN. In *ACM WWW*, Edinburgh, UK, May 2006.