Client perceived performance in a campus network with a wireless LAN controller

Guillaume-Jean Herbiet, *Master student, College of Computing*
Sriram Lakshmanan, *PhD student, School of Electrical and Computer Engineering*

Abstract— **Wireless Local Area Network Controller Systems (WCS) are recently introduced platforms for wireless LAN planning, configuration, and management. Such systems eliminate the need for manual configuration of Access points and also provide several advantages such as dynamic power control and channel selection for better quality of service, better interference management, improved security services and network management. Although there are several commercial products available, a proper study of the actual end-user benefits of such a technology has not been performed. In this work, we perform extensive measurements in a campus network in which a WCS controls more than 60 access points (APs). We specifically focus on mobility management and quality of service features of WCS. The experimental results show that the use of WCS doesn't really improve the latency of the complete roaming process, which clearly contrasts with industry claim. However, our sets of measurements showed that the consistent throughput available is improved and that experienced loss rate at the transport and MAC layers is much lesser when the wireless architecture relies on a central controller. Those metrics tend to indicate that introduction of WCS can help in making wireless networks more robust and achieve better performance with mobile applications.**

I. INTRODUCTION

Wireless Local Area Networks have proliferated significantly with widespread dense deployments in campuses, hotspots, airports and even homes. The deployments have reached very large proportions currently. Enterprises and organizations have felt the need for wireless network management. A recent development in this direction is WCS.

Wireless LAN Controllers Systems (WCS) are platforms for wireless LAN planning, configuration, and management. They provide a foundation that allows IT managers to design, control, and monitor enterprise wireless networks from a centralized location offering tools for wireless LAN planning and design, RF management, location tracking, Intrusion Prevention System (IPS), and wireless LAN systems configuration, monitoring, and management.

While WCS primarily provide a central control point for enterprise WLANs enabling ease of management and security enforcement, they also reduce the complexity of the access points. Further, the controller is empowered to take globally optimal decisions as opposed to a local optimal decisions that conventional WLAN architectures allow. Also, dynamic channel assignment and power control can potentially reduce interference and contention among users. Thus WCS have significant performance benefits and management benefits.

Although the technology sounds very promising, no extensive study has been made so as to assess the actual improvements in terms of user performance, particularly regarding quality of service and management of the mobility of the end user. Such a study is valuable not just as an academic exercise but would provide valuable information for network managers, helping them to make calculated decisions in future deployments, for wireless application developers, who require network-level metric assessments so as to design more efficient products, and for manufacturers of controllers for refining their algorithms and identifying problems with their current solutions.

In this work, we focus on evaluating the client performance and mobility benefits that WCS provides. Particularly we analyze the latency of the roaming process from an access point to another under several scenarios of pedestrian mobility on networks with or without a WCS. We also identify to what extent the throughput is affected by the handoff process. A third important metric we observed is how the loss rate varies due to the better RF management performed by the controller. Our purpose is to see to what extent the introduction of WCS may improve the use of different applications in a mobile wireless context from an end-user perspective.

We resort to extensive measurements in the Klaus Advanced Computing Building on the Georgia Tech campus where the WCS is used to control access to the wireless network and in other buildings of the campus, with similar AP density and usage profile, where such an architecture is not implemented.

In performing those measurements, we also identify the benefits that current WCS provides in terms of throughput improvement and wireless losses but also their shortcomings when roaming improvement is concerned. To our knowledge, this is the first step in experimental analysis of WCS performance.

We observed that the overall handoff latency does not improve significantly with the WCS. Throughput and loss rates are improved about 20% on the average. Besides, performance varies with the card used and thus the controller benefits are not irrespective of the card used.

The remainder of this report is organized as follows. Section II described related work. Section III provides the background on WCS and the benefits envisioned. The measurement methodology for the study is presented in Section IV. Section V details results of our analysis, while section VI provides a conclusion on the WCS performance and section VII introduces future work.

II. RELATED WORK

Although there is very less academic literature on wireless controllers themselves, studying wireless campus networks and latency during roaming have been popular topics in networking.

A. Studies on campus wireless networks

In [3-7,17], the authors study and analyze various aspects of campus wireless networks such as user association pattern, applications used and other network performance metrics using traces generated through actual measurements. However, none

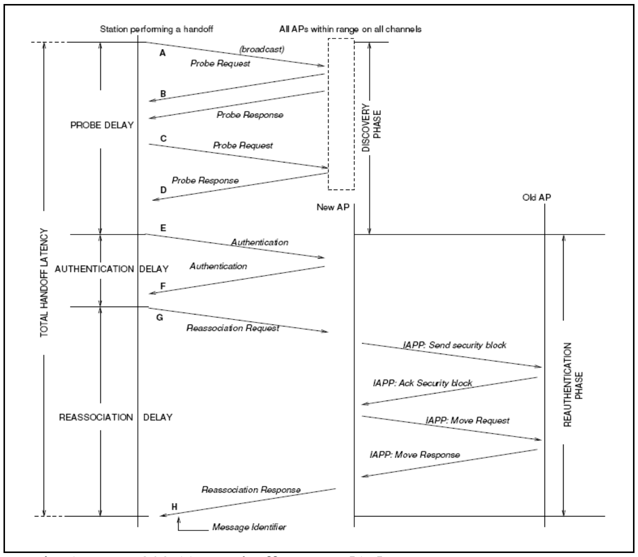

Fig. 1: IEEE 802.11 Hand-off process [17]

of these studies are carried out with a controller system in place. Thus, the results and trends are not directly applicable to the scenario considered in this work but help us understand the particular implantation and activity profile of wireless campus networks.

In [3], Kotz and Essien study patterns of activity and mobility in a large campus network. Their work reckon that a wide variety of traffic from hour to hour, day to day, week to week, with regular patterns, different from those common in enterprise WLANs. Besides, they notice that many wireless cards are aggressive in association, leading to a significant number (about 17%) of short sessions involving roaming, even if the user remain stationary (40% of session with roaming).

In [5], Henderson and al., conduct a similar study in the same network several years after and revealed significant changes. First, mobility and thus roaming between AP's increased. Besides, development of always-on devices cause the number of users associated with individual AP to soar. Furthermore, wireless VoIP and peer-to-peer traffic appeared. Calls and large file transfers may severely impact other users, mainly due to wireless channel access specification, and make wireless-specific management, like the introduction of VLANS with different QoS levels, desirable.

Another interesting work was realized by Jain and al. in [17], when they attempted to model user registration patterns to several access points. Their results reckon that user registration showed a hierarchy and that AP's can be grouped in clusters based on the probability of inter-AP transitions. Besides the sizes of those clusters are highly skewed. The authors therefore plead for protocols oriented towards the election of cluster heads that will perform local control and signaling functions.

By reviewing all those studies, it appears that the different functions performed by a wireless controller (such as QoS and mobility management) are likely to enhance end-user performance in today's wireless campus networks.
Another important point, detailed further in the methodology section, concerns the existence of traffic patterns., and made us give a consideration to the particular times where our measurements where realized.

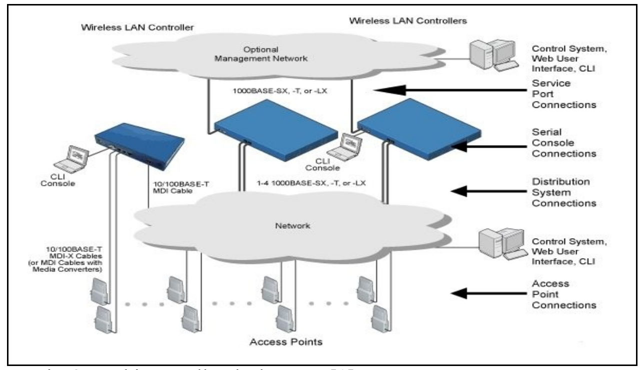

Fig. 2: Multi controller deployment [1]

B.Hand off latency in wireless LANs

In [8], the authors analyze the handoff process and describe the main contributions of the handoff latency. Particularly, they observe that handoff latency could be large (i.e from 120ms to 500 ms) and varying. The dominant component of handoff latency is the probe function. [14] describes handoff measurement techniques and also provides a detailed description of the handoff process. [9-13] are solutions that are designed to minimize the handoff latency. However, the solutions used are developed for the conventional thick APs and each solution has different requirements at the clients and the APs. All these studies, while relevant, do not answer the questions raised in the current work.

While there exist several commercial products which describe a controller based network management solution [1, 15, 16], the details of the algorithms used are proprietary and hence not available.

III. WIRELESS NETWORK CONTROLLER

Wireless Controller solutions are designed to simplify deploying and managing large-scale wireless LANs. The operating system manages all data client, communications, and system administration functions, performs Radio Resource Management (RRM) functions and manages system-wide mobility policies.

A. Principles and deployment

A typical WCS solution architecture consists of one or several wireless controllers and their associated lightweight access points controlled by the operating system. All devices can be controlled through a web-based, command line or proprietary interface.

B.Claimed advantages

The following paragraphs describe some key points on which wireless controllers claim to yield improvements. However, we focus only on network performance related benefits and do not consider other advantages such as security and network management.

1)Quality of service

Most WCS systems support QoS by assigning a specific level of service to each VLAN created on the network. So as to achieve this, they rely on specific IEEE 802.11e user profiles for over-the-air traffic. The access points then use this profile to derive IP DSCP values that are visible on the wired LAN.

2)Radio resource management

The radio resource management (RRM) software embedded in the controller acts as a built-in RF engineer to consistently provide real-time RF management of the wireless network. Processes are separated from 802.11a and 802.11b/g. RRM enables controllers to continually monitor their associated lightweight access points for metrics like traffic load, interference (amount of traffic for other 802.11 sources), noise (amount of traffic from non-802.11 sources), coverage (signal strength and signal-to-noise ratio from connected clients) and number of nearby APs.

Using this information, RRM can periodically reconfigure the 802.11 RF network for best efficiency. To do this, RRM performs functions like dynamic channel assignment (changing channel used by nearby APs), dynamic transmit power control, coverage hole detection and correction, client and network load balancing (by automatically forcing some subscribers to associate with nearby access points, allowing higher throughput for all clients).

3)Mobility management

In multiple-controllers deployments, WCS defines three types of client roaming, that are managed differently.

- *Same-subnet (layer 2) roaming* concerns roaming between two APs managed by the same controller. and is transparent to the client as the session is sustained and the client continues using the same DHCP-assigned or client-assigned IP Address. The controller provides DHCP functionality with a relay function. On the Georgia Tech Local Area Wireless Network (LAWN), all users belong to the same subnet and all thin access points are controlled by a single controller. So this is the roaming process that we studied in this work.

- *Inter-controller (layer 2) roaming* allows to move between APs managed by different controllers in the same subnet and also transparent to the client, as the session is sustained and a tunnel between controllers allows the client to continue using the same DHCP- or client-assigned IP Address as long as the session remains active.

- *Inter-subnet (layer 3) roaming* supports client roaming across access points managed by controllers in the same mobility group on different subnets. This roaming is transparent to the client, because the session is sustained and a tunnel between the controllers allows the client to continue using the same DHCP-assigned or client-assigned IP Address as long as the session remains active.

These techniques allow, according to wireless controllers manufacturers, to significantly reduce the latency and throughput degradation while roaming from an access point to another and are the key point of our study.

IV.METHODOLOGY

This section describes our experiment setup and the methodology used for measurement, the metrics we observed in our experiment, with all the assumptions and limitations related to the fact that we are operating in a real environment where some parameters are hardly controllable.

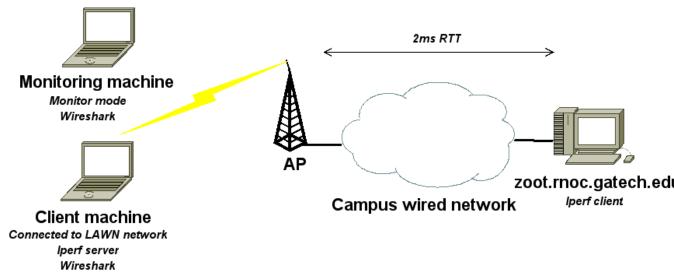

Fig. 3: Experiment setup

A.Experiment setup

The experiments were carried out using two laptops. The first one, a HP Pavillion dv1160ea laptop running Ubuntu Linux and using an 802.11b/g Intel card, was used in monitor mode so as to capture control frames of the IEEE 802.11 that will help identify the roaming process. The second laptop, a Dell Inspiron E1405 with Windows Vista and using a DELL 802.11g card , acted as a regular client machine of the campus wireless network while moving at pedestrian speed.

The client machine was simulating a TCP or UDP traffic while running iperf with a host as close as possible from the measurement network and as lightly loaded as possible (namely zoot.rnoc.gatech.edu). On both machines, Wireshark was used to capture and pre-process the frames and packets. Those traces were further analyzed to provide the results described in the next section.

B.Measurement conditions

Our measurements consisted of ten-minute rounds across a building, while taking several pre-defined paths under pedestrian mobility. During capture, we recorded the places where hand-off occurred and the APs which the client was roaming from and to so as to have an estimation of the optimality of the associations.

The experiments were conducted at different times of the day and different moments of the week, on week-days and week-ends, so as to capture the behavior under different ambient load. To identify the traffic patterns (number of users connected, amount of data transferred), we referred to different measurement tools managed by the Office of Information Technology, which includes lists of access points in function in different buildings, average coverage level and logs from the MRTG tool that constantly monitor the traffic from and to the different access point in the Georgia Tech wireless network.

As on the Georgia Tech campus, only the Klaus Advanced Computing Building is currently implementing a wireless architecture based on controllers, measurements were performed there first. So as to obtain comparison points, we performed similar measurements in other buildings on the campus so as to gather results from WCS-based and “thick”-based networks. We tried to ensure that those different buildings had a similar AP density, based on analysis of building structure and access-points implantation and raw measures of reception quality. The selected buildings are College of Computing building, Student center and instructional center.

C.Metrics used

From the results of the experiments we first seek to assess the end-user perceived performance in a WCS-based wireless architecture, in terms of roaming latency, achieved throughput and loss rate. Then those results are put in perspective with similar assessment performed in traditional thick-AP architecture. Therefore we focused on the following metrics.

For mobility management, the monitoring machine recorded the control frames sent and received by the client host. We then processed this capture so as to compute the overall time spent in probing, roaming and sending actual data. We also extracted the different hand-offs that occurred during the recording, with the associated probing and reauthentication delays.

During the measurements, the client machine recorded the TCP or UDP traffic it received from the remote wired node. The captured packets were then processed to compute the evolution of the throughput. The TCP throughput over 1 second intervals was calculated from the packet capture.

Finally, to obtain an estimate of losses , retransmissions were calculated at both the MAC layer and TCP layer by appropriate filtering of the packet capture. The MAC layer retransmissions are indicative of the actual wireless losses and are counted from those packets which have the 'retry' bit set to 1. On the other hand, TCP layer retransmissions are indicative of the congestion related losses.

D.Assumptions/Considerations

Our work aims to be a in-field set of experiments, making harder to control external factors than it would be if we were to use a lab setting. However, we tried to control some of those parameters and worked under the following assumptions.

1)Wireless bottleneck

The experiments are intended to identify the wireless performance. However, the Iperf source is located on a wired machine. Thus it is important to ensure that the measured performance is due to the wireless bottleneck rather than the wired bottleneck. Therefore, a machine with a small RTT of 2ms was chosen within the Georgia Tech campus.

The premise that the wireless is the bottleneck is confirmed by the fact that the static rates upto 23 Mbps was observed. Further, the rate was varying much, indicating the load on the wireless leg.

Although, the available bandwidth was tested before experiments, any increase in wired traffic which could shift the bottleneck is not accounted for in our experiments. However, we believe that such an event is less likely in the test environment.

2)TCP Window setting in Iperf

When using TCP as the transport protocol to measure available throughput, it is important to set the parameters of the TCP connection appropriately. The TCP window size is set to a higher value than the bandwidth delay product of the path, so that TCP source will not be window limited. The Bandwidth delay product becomes: *RTT*BW= 2ms*54 Mbps = 108 Kb =13.5 KB*. Hence a setting of *16 KB* was used.

3)Perfect capture

In our processing of the files, we assumed that all management packets received by client are captured by sniffer and all packets received on interface captured by Wireshark.

The packet capture tool, Wireshark provides the number of packets dropped while performing the capture. This was monitored for each run and the number of packets dropped in Wireshark = 0 for every run, except one for which 326 of 1013001 packets dropped. Even in this case, the fraction was very low.

4)Coverage ensured by ping

Before performing the experiments, the coverage was ensured using a train of PING. The places of no coverage were filtered from the measurement.

5)Load on APs do not vary drastically

The throughput results would be skewed is the load at the APs with WCS and those without WCS. Since the client induced load was not under our control, the test locations and times of the day were chosen so that the usage pattern was similar between the building with WCS and without WCS. However, the results presented here do not explicitly account for load variations.

V.Results and Analysis

This section presents the results from our different measurement sessions. Recalling the different assumptions and limitations stated in the paragraph above, a detailed analysis is then presented for each metric of interest.

A.Hand-off performances

The results concerning the hand-off performances experienced with a WCS system will first be presented and then compared to results obtained in building where such a controller is not present.

1)Performances with a controller

The series of experiments realized in the Klaus building where the WCS is presents allowed us to record a total of 41 roamings during 10 different measurement sessions. The management frames captured by the monitoring node were then exported and post processed so as identify the total length of the recording, the total length of data exchange (and its ratio to the experiment length), the total length of all the probing and reauthentication phases (and their ratio to the experiment length). Besides, each roaming was identified based on the authentication and reassociation frames and was associated with the preceding probing period. The repartition of the different states and the data regarding the roaming processes are summarized in the tables 1 and 2 respectively.

The first analysis of our results sets showed some points where the probing phase was exceeding several seconds. Most of them are located at the beginning of an experiment and correspond to the first probing period, right after the client node has turned its wireless interface up. As they didn't correspond to a mobility hand-off, they were removed from our results. Some other points also showed probing phases with the same order of length. As they were mainly composed of probe requests and few probe responses, we attributed those periods to connectivity loss and also removed them from our set. In the remaining of this study, the corrected set is analyzed.

After those corrections, the roaming durations ranged from 2 ms to more than 600 ms with a average around 300 ms. Those

Table 1: Mobility phases repartition for the Klaus experiments

	Capture (s)	Data (s)	Data ratio	Probing (s)	Probing ratio	Authentication (s)	Reauthentication (s)	Rauthentication ratio	Number of roamings
TOTAL	4928.83	2876.14	58.35 %	1902.24	38.59 %	1.15	1903.39	38.61 %	41
AVERAGE	492.88	287.61	58.35 %	190.22	38.59 %	0.11	190.33	38.61 %	4.1
STDEV	228.79	153.37	67.03 %	105.74	46.21 %	0.14	105.88	46.28 %	3
MAX	917.19	555.78	60.59 %	358.40	39.07 %	0.43	358.83	39.12 %	11
MIN	217.04	95.38	43.94 %	0.55	0.25 %	0.02	0.56	0.26 %	2

Table 2: Details of roaming latency for Klaus experiments

	Roaming	Probing	Probing Ratio	Total
AVERAGE	0.033 s	0.260 s	86.567 %	0.293 s
STDEV	0.068 s	0.137 s	21.141 %	0.165 s
MAX	0.328 s	0.562 s	99.547 %	0.634 s
MIN	0.002 s	0.000 s	0.000 %	0.002 s
MEDIAN	0.022 s	0.284 s	92.777 %	0.303 s

values are close to typical hand-off latencies achieved by thick-AP infrastructure[8] and way below the 2 ms claimed by the manufacturer of the wireless controller we tested[1].

A noticeable aspect of those results is the predominance of the probing phase during the hand-off process. Due to the centralized management of mobility performed by the wireless controller, we expected that this phase would be reduced, as the WCS could, for instance, let only the most suitable AP answer to the probe requests send by the client prior to roaming. Here this phase still represents more than 85% of the hand-off process.

Another point we wanted to investigate is the impact of the ambient load on mobility performance. For this reason we performed experiments during rush periods (on week-days during office hours) and on week-ends. However, variation of the mean roaming latency is very low from one session to another, even when variation in the ambient load is shown in the MRTG log from LAWN monitoring. Our insight is that, as Klaus building has a high AP density, the wireless network there is always used below a saturation point where mobility performance can be dramatically affected.

2)Comparison with traditional AP architecture

So as to compare the previous results to cases where mobility occurs with a traditional wireless architecture, we performed experiments in other buildings on the campus that don't use the WCS. A total of 21 roamings in five different sessions were recorded under different ambient load, trying to reproduce similar conditions as for the different Klaus records.

Here again some points were removed, for the same reasons detailed before, and our measurement set narrowed down to 15 measurement points. Statistical analysis of those results are showed in the tables 3.

Those results show that the probing phase with a thick-AP based architecture is more important than when a wireless controller is involved. This could lead us to conclude that the

Table 3: Details of roaming latency for thick APs experiments

	Roaming	Probing	Probing Ratio	Total
AVERAGE	0.004 s	0.279 s	92.130 %	0.283 s
STDEV	0.005 s	0.133 s	24.712 %	0.132 s
MAX	0.022 s	0.403 s	99.556 %	0.406 s
MIN	0.002 s	0.000 s	0.000 %	0.002 s
MEDIAN	0.002 s	0.338 s	99.325 %	0.340 s

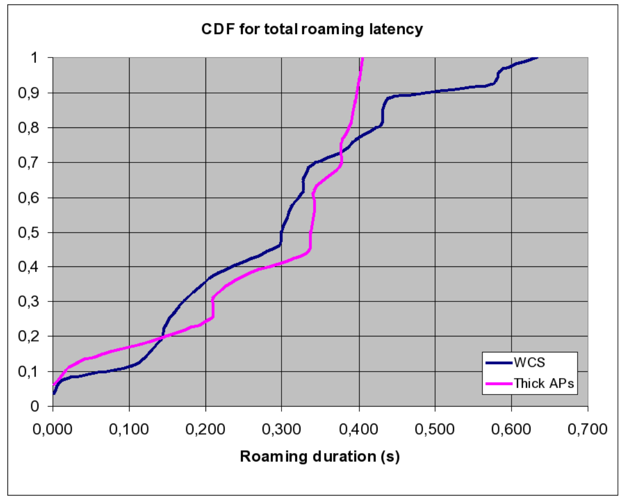

Fig. 4: CDF for total roaming latency with and without WCS

WCS still lead to some improvements during the probing phase. However, as our experiments were conducted in different building, those results can be attributed to differences in the access point implantation, even if the trend described here was verified in all the buildings where we run the comparative experiments.

Besides this remark, the performances achieved for roaming latency are very similar to what was experienced in presence of the controller. As showed on the figure 4, the distribution of roaming latency in the Klaus building is broader than in thick-AP architectures, even if the median value is smaller with the WCS.

However we tried to perform as much experiments as possible, it is impossible to ensure that those differences will not vary if another data set was to be collected. Besides, the numbers are so close that a mainstream end-user is very unlikely to perceive noticeable difference while moving from an access-point to another, whatever the architecture in use.

B. Throughput analysis

In this section, the result of the TCP throughput achieved with pedestrian mobility are presented. Specifically, the average , min,max throughput achievable are shown in the table 4.

As can be observed in the figure, the average throughput is higher with the presence of WCS than without. However, the minimum and maximum values are similar. Also, it must be noted that , while the static throughputs obtainable are as high as 23 Mbps in the WCS case, this reduction in the maximum throughput is only due to mobility.

The figure 5 shows the CDF plot of the instantaneous throughput taken over 1 second intervals. Again, it can be observed that the WCS case shows a better throughput.

Table 4: Throughput results with and without a controller system

Throughput	WCS	Thick APs
AVERAGE	9.37 Mbps	7.73 Mbps
MIN	10.9 Kbps	32.3 Kbps
MAX	19.2 Mbps	18.9 Mbps

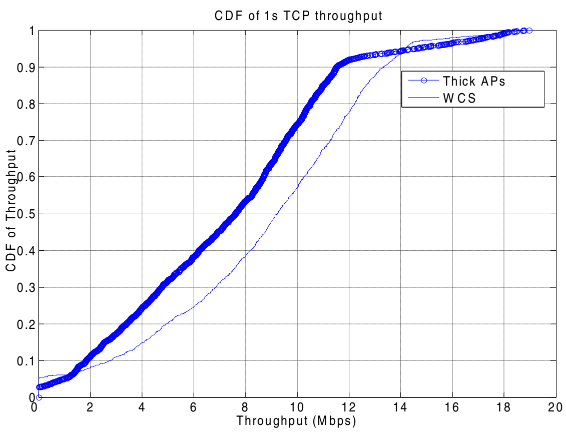

Fig. 5: CDF of TCP throughput with and without a WCS system

Table 5: Individual TCP Throughput statistics

Configuration	Average (bps)	Minimum (bps)	Maximum (bps)
K1	9904215	1411776	17477568
K2	8708971	54720	19086336
K3	10077273	1805760	14533632
K4	9915142	10944	18407808
K5	8277287	875520	15715584
			17044185
C1	6047983	481536	11086272
C2	7958249	32832	18965952
C3	6901655	109440	18484416
C4	8426301	733248	18035712
C5	9353558	1926144	14686848

The table 5 of individual throughput results is shown above. The measurements in the WCS setting are given by K1 to K5, whereas the measurements in the Thick AP setting are given by C1 to C5.

C.Loss rate

Retransmissions occur at both the MAC layer and Transport layer due to wireless related losses. Random wireless channel losses, could create actual packet errors and/or buffer drops at the sender due to the retransmissions. Hence, the number of MAC layer retransmissions is a representative measure for the loss rate.

1)MAC layer retransmissions

The MAC layer retransmission statistics for the WCS and thick AP cases are shown in the tables 6 and 7.

In the tables, Klaus1 t oKlaus6 represents the measurement experiments in the Klaus building equipped with a WCS. Since a flow was setup to the wireless client, the upstream and downstream statistics are shown as 'up' and 'down' respectively. Similarly, for the experiments in the other buildings equipped with thick APs, CoC1,CoC2 represent the measurements in the CoC building. Similarly, IC and SC1,SC2 represents the measurements in the Instruction center and Student center buildings. The CoC2, value shows entries of zero due to a problem with the packet capture.

Table 6: MAC layer retransmissions with WCS

	Data Retx	Data Recd	Man. Retx	Man Recd	Data loss Rate %	Man. Loss rate %
Klaus1 (up)	12733	51548	3	400	19.8	0.75
Klaus1 (down)	16440	83066	145	295	19.8	49
Klaus2 (up)	18406	68035	0	1001	27.0	0
Klaus2 (down)	24025	97935	311	625	19.6	33
Klaus3 (up)	3016	7081	1	328	29.8	0.3
Klaus3 (down)	2944	10827	114	243	21.3	31.9
Klaus4 (up)	3227	11272	2	616	22.2	0.3
Klaus4 (down)	3146	10519	150	343	23.0	30.4
Klaus5 (up)	2617	7364	1	380	26.2	0.2
Klaus5 (down)	2305	8026	128	261	22.3	32.9
Klaus6 (up)	2317	8258	0	381	21.9	0
Klaus6 (down)	2159	6039	89	206	26.3	30.1
	93335	369970	944	5079	**20.1**	**15.6**

Table 7: MAC layer retransmissions for thick-AP architecture

	Data Retx	Data Recd	Man. Retx	Man Recd	Data loss Rate	Man. Loss rate
Coc1(up)	1847	3461	0	152	34.7	0
Coc1(down)	1509	4853	166	370	23.7	30.9
Coc2(up)	0	0	0	161	0	0
Coc2(down)	0	0	164	341	0	32.4
IC(up)	517	1660	0	284	23.7	0
IC(down)	151	245	247	878	38.1	21.9
SC1(up)	10198	22598	0	669	31.9	0
SC1(down)	3592	18453	165	355	16.2	31.7
SC2(up)	3743	7875	1	545	32.2	0.1
Sc2(down)	1642	9983	122	274	14.1	30.8
	23199	69128	865	4029	**25.1**	**17.7**

Overall, one can observe the following main points. The loss rates are different between the WCS and Thick AP case. More specifically, while WCS has a 20% retransmission, Thick APs have a 25% retransmission. This indicates that a 20% improvement in retransmission is observed with the WCS. There could be several reasons for this effect, like better interference management, cleaner environment, more powerful and intelligent APs; etc. However, on the whole the newer WCS system (the algorithms and the equipment) together contribute to an increase in throughput and a consequent decrease in retransmissions.

The management and ACK packets (down) are smaller in size. The management packet definitely has a lesser loss rate due to the smaller size. This can be explained using the dependence of loss rate on packet size , for a fixed Bit error rate.

Table 8: Individual TCP Throughput statistics

TCP Effects			
Location	Retransmissions	Successful	Fraction (%)
WCS	529	2086483	0.025
Thick APs	1021	2088461	0.0488

There is also an asymmetry in the upstream and downstream loss rates.

2)TCP level retransmissions

TCP layer retransmissions can be caused due to buffer drops and timeouts induced by wireless channel impairments. This effect was also monitored.

Overall, we observe that the TCP level retransmissions in the WCS environment is 50% lesser than that in the Thick AP environment. This can be attributed to lesser congestion perceived by the transport layer. Better channel contention and interference management could lead to lesser buffer drops and timeouts, contributing to lesser retransmissions.

D.Other findings

In this paragraph, we would like to mention some findings we made during our series of experiments. As those remarks were not the main purpose of our study, only limited analysis on them can be performed and they are mentioned as interesting directions for further work.

1)Importance of client card

We extended our measurement sessions in the Klaus building where a WCS architecture is implemented, by introducing different wireless cards for the client host, using both the 802.11b and 802.11g standard. Namely, the different cards used were DELL 802.11g, Intel 802.11 g, Linksys 802.11g, Linksys 802.11g - SRX. The insight of this experiment was to assess the importance of the wireless card roaming algorithm in presence of a wireless controller. The results concerning the roaming latencies experienced for two ot the cards are summarized in the tables 9 and 10.

Those numbers show a different behavior for the cards regarding the absolute duration of the roaming process and the relative extent of the probing phase, even if the experiments were realized in the same building, taking the same paths and under similar background load.

This observation is important and should be stated in a study assessing the achieved performances from a user point of view. As end-users may use a variety of different chipsets for their wireless cards, they may behave differently, depending on the hand-off algorithm they implement. However, the results have same order of magnitude and seem to indicate that what the user will perceive as performances won't be too different from one card to another.

2)Association optimality

As stated in the measurement conditions paragraph, during our measurements in the Klaus building, we recorded the current position of the host (as nearest room number) along with the identifier of the APs the node was roaming from and to. As the access point on the Georgia Tech campus are identified by their building number and the room they are in or

Table 9: Details of roaming latency for Intel 802.11 g

	Roaming	Probing	Probing Ratio	Total
AVERAGE	0.019 s	0.433 s	65.609 %	0.452 s
STDEV	0.012 s	0.746 s	41.302 %	0.747 s
MAX	0.040 s	2.071 s	98.760 %	2.097 s
MIN	0.007 s	0.002 s	4.412 %	0.025 s
MEDIAN	0.02 s	0.17 s	95.52 %	0.18 s

Table 10: Details of roaming latency for card Linksys 802.11g

	Roaming	Probing	Probing Ratio	Total
AVERAGE	0.032 s	0.386 s	87.434 %	0.418 s
STDEV	0.065 s	0.514 s	20.591 %	0.519 s
MAX	0.328 s	2.698 s	99.547 %	2.720 s
MIN	0.002 s	0.000 s	0.000 %	0.002 s
MEDIAN	0.02 s	0.28 s	92.88 %	0.3 s

closest to, this allowed us to quickly evaluate the closeness of the new AP that was selected for reauthentication by the client card.

This recording, first designed to be added as metadata to our results, didn't rely in accurate positioning (using for instance a GPS device) and isn't correlated with record of metrics like signal to noise ratio at the client node. So identification of roaming point is not completely accurate, and the notion of “optimal” access point to roam to is only based on a estimation of the distance between the client node and the access point.

However, our measurements records clearly show that the client node is not always associated to the optimal access point. This particularly happens after the client node has been moving, and associated to an access point that was optimal at that time, and new stands still, closer to a new access point. As most wireless card use an hysteresis algorithms, the node remains connected to the suboptimal AP.

We though that the WCS having a complete picture of AP implantation, and the access-points reporting the signal to noise ratio they experience with each client, this architecture can be used to solve this problem, for instance by forcing deauthentication of the client at the suboptimal AP.

This aspect is definitely worth a dedicated set of measurements and a great open issue for further work.

3)Impact of extended-range cards

The SRX cards use multiple antenna technology to provide range extension. The card was thus able to identify and receive successfully the beacons of more APs than other cards. Moreover, the problem of suboptimal association was exacerbated with this card , since the card continued to remain associated with a far-away AP. For instance, when the experiments were started at a corner in the third floor of the Klaus building, the associated AP belonged to the second floor. The AP ids for each floor was identified using the APlist data provided by Dr.Clark. The main observation is that when combined with Auto Rate Fallback, the persistence of AP association leads to significant throughput degradation. Thus , while the signal strength decreased when moving away from the associated AP, the Autorate fall back technique caused a decreased rate before initiating a handoff. Thus, the persistence aspect of mobility management at the client causes more severe problems with range extension cards. This is an important observation given that the cards are based on draft 802.11n standard and the standard is expected to be finalized soon.

VI. Conclusion

The WCS doesn't seem to noticeably improve the mobility management, as it is mostly managed by the client wireless card and not the controller.

However, throughput and loss performance were better when a WCS is used, which could be interpreted as a benefit from the controller or be due to the particular AP implantation in the Klaus building. Thus, the new APs deployed in the Klaus building along with the controller algorithms, together give higher throughputs and lesser losses.

VII. Future Work

The paragraphs presenting some side conclusions of our work also introduced some interesting ideas so as to complete our work. Particularly, the problem of the optimality of the selected access-point and the eventual impact of the wireless LAN controller on this selection is really worth being investigated as it directly impacts the end-user performance for its communication.

As our work wanted to assess what a user will perceive on his daily usage of the network, we mainly focused on field experiments where all side parameters couldn't be precisely controlled, despite the efforts we put in selecting the buildings, paths and moments to perform the experiments. Therefore, we would like to complete our set of measurements with lab experiments where we would be able to investigate more thoroughly various effects of the introduction of the wireless controller, like the direct impact of RF management in the achieved throughput and the rate of losses.

Finally, testing the WCS behavior on inter-controller and inter-subnet roaming situations, which are not present in the Georgia Tech wireless network, is also worth of interest.

Acknowledgment

The authors would like to thank Dr. Russel J. Clark and Earl Barfield, from Office of Information Technology, for their constant help throughout the project.

References

[1] Cisco documentation on Wireless Control Systems, http://www.cisco.com/en/US/products/ps6305/

[2] A. Balachandran, G.M. Voelker, P. Bahl, and V. Rangan, "Characterizing User Behavior and Network Performance in a Public Wireless LAN," *Proc. ACM SIGMETRICS*, June 2002.

[3] D. Kotz and K. Essien, "Analysis of a Campus-Wide Wireless Network," *Proc. ACM MobiCom '02*, Sept. 2002.

[4] D. Schwab and R. Bunt, "Characterizing the Use of a Campus Wireless Network," Proc. IEEE INFOCOM, Mar. 2004.

[5] Tristan Henderson, David Kotz, and Ilya Abyzov "The changing usage of a mature campus-wide wireless network" *Proc. ACM MobiCom*, pages 187-201, Philadelphia, PA, USA, September 2004.

[6] Wei-Jen Hsu and Ahmed Helmy. On modeling user associations in wireless lan traces on university campuses**.** *Proceedings of the Second Workshop on Wireless Network Measurements (WiNMee 2006)*, Boston, MA, USA, April 2006.

[7] Francisco Chinchilla, Mark Lindsey, and Maria Papadopouli. Analysis of wireless information locality and association patterns in a campus. *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 906-917, Hong Kong, China, March 2004. IEEE.

[8] A. Mishra, M. Shin, and W. Arbaugh, "An Empirical Analysis of the IEEE 802.11 MAC layer Handoff Process" *ACM Computer Communications Review, vol. 33, no. 2, Apr. 2003.*

[9] M. Shin, A. Mishra and W.A. Arbaugh, "Improving the Latency of 802.11 Hand-offs using Neighbor Graphs" *Mobisys 2004 June, 2004, Boston, USA.*

[10] I. Ramani and S. Savage, "SyncScan: Practical Fast Handoff for 802.11 Infrastructure Networks" *Proceedings of the IEEE Infocom, March 2005.*

[11] F. K. Al-Bin-Ali, P. Boddupalli, and N. Davies, "An Inter-Access Point Handoff mechanism for Wireless Network Management: The Sabino System" *in Proceedings of the International Conference on Wireless Networks , Las Vegas, NV, June 2003*

[12] Vladimir Brik, Arunesh Mishra, Suman Banerjee "Eliminating handoff latencies in 802.11 WLANs using Multiple Radios: Applications, Experience, and Evaluation ", *IMC 2005*

[13] Arunesh Mishra, Vladimir Brik, Suman Banerjee, Aravind, Srinivasan, William Arbaugh "A Client-driven Approach for Channel Management in Wireless LANs.". *IEEE Infocom*, Barcelona, Spain, 2006.

[14] Francisco Gonzalez, Jesïs Pιrez Dvaz, Victor H. Zarat e, "HAMS: Layer 2 Handoff Accurate Measurement Strategy in WLANs 802.11 " *WinMee 2005.*

[15] Meru Networks,www.merunetworks.com

[16] Aruba Networks , www.arubanetworks.com

[17] R. Jain, D. Lelescu; M. Balakrishnan, "Model T: An Empirical Model for User Registration Patterns in a Campus Wireless LAN", *Mobicom 2005.*