SmartAP: Practical WLAN Backhaul Aggregation

Eduard Goma
Telefonica Research
Barcelona, Spain
Email: goma@tid.es

Domenico Giustiniano
ETH Zurich
Zurich, Switzerland
Email: domenico.giustiniano@tik.ee.ethz.ch

***Abstract*— In light of the growing disparity between residential broadband and 802.11 speeds, Access Point (AP) backhaul aggregation has been proposed by the research community as a service whereby residential customers may enjoy higher throughput when in range of participating 802.11 APs. The fundamental assumption of these works is that 802.11 clients can be modified at driver level. However, due to the high diversity of 802.11 chipsets and drivers in the market, the cost of modifying any WiFi driver for any operating system is prohibitive, which in turn makes the solution unpractical for commercial deployment. In this paper, we propose a WLAN backhaul aggregation scheme that works with unmodified 802.11 clients. We introduce SmartAP, a single-radio 802.11 AP, that can deliver higher throughput to its off-the-shelf clients by aggregating the backhaul capacity of APs in range, *even if these APs are in different radio channels*. SmartAP reaches this goal without adding new radio hardware to the network. We build SmartAP with off-the-shelf hardware and evaluate its performance in a network testbed of 6-nodes with unmodified smartphones and laptops as clients. We evaluate the conditions under which SmartAP leads to gains comparable to state-of-the-art approaches with client-side modifications and we demonstrate that using a two-hop transmission to access additional backhaul capacity yields substantial benefits.**

I. Introduction

Over the last decade, Internet access speeds have significantly increased due to new network deployments such as xDSL, Fiber To The Home (FTTH) or cable. Out of these technologies, the most widespread is xDSL that can offer theoretical speeds of up to 50 Mb/s to users that are close to the DSLAM. However, the reality is that the Internet access speeds remain low in many regions of the world. For instance, 97% of the residential access speeds in Colombia in 2011 were below or equal to 4 Mb/s [1]. As a way to address this problem, wireless backhaul bandwidth aggregation has been proposed as a means to increase users' connectivity speed by exploiting the higher speed of 802.11 WLAN with respect to the Access Point (AP) backhaul [2], [3]. In addition, commercial initiatives like FON [4] aim to create nation-wide WiFi sharing communities that offer the possibility for users to connect to the Internet whenever they are in range of an AP from the community. Internet Service Providers (ISPs) running wired and cellular access networks see backhaul bandwidth aggregation as a potential approach to increase the experienced throughput without major short-term investments to their backhaul networks and WiFi sharing as an opportunity to offload traffic from the cellular network.

While WiFi sharing is a reality in today's networks, the opposite holds for backhaul bandwidth aggregation methods. Taking an academic concept to a commercial deployment brings up a number of factors that need to be considered. First, the overall solution costs must be kept low. Thus, introducing extra hardware or requiring highly specialized devices is not a commercial option. Second, the deployment of the solution must not be too cumbersome, otherwise the approach would not scale with the number of devices deployed. Current backhaul bandwidth aggregation solutions [2], [3] fail to meet the requirements above. The reason is that they rely upon the assumptions that the wireless of every possible WiFi client, including smartphones, tablets, laptops, etc., can be modified, which is clearly cost prohibitive. The problem of the diversity of devices that need to be modified can be solved by changing the aggregation point in the system. Our idea is to deploy the aggregation scheme in a much smaller set of devices, the APs, which are usually provided by ISPs. Introducing a secondary WiFi radio in the APs could be a technical solution, but it would increase the cost of a device that is subsidized by the ISP, making the solution impracticable.

Our approach is to design and develop a solution that can enable the desired functionality simply through *software AP modifications* and *without any client support*. In this paper we present SmartAP, a single-radio AP that allows ISPs to commercially deploy wireless backhaul aggregation, simply with a *remote firmware upgrade* to their installed APs base. SmartAP aggregates the unused capacity of neighboring APs, regardless of their radio channel and offers it to off-the-shelf 802.11 clients in range.

Our contributions are multi-fold and they can be summarized as follows:

1) we propose a feasible and cost-effective backhaul bandwidth aggregation solution.
2) we formulate the problem of wireless backhaul bandwidth aggregation through a single-radio AP and show that the problem can be mapped to the client-based solutions allowing the same optimization objectives.
3) we evaluate the performance trade-off of the AP- and client-based aggregation schemes both analytically and experimentally.
4) we implement SmartAP and show through experiments that it is able to aggregate the bandwidth of neighboring backhauls and increase the throughput of off-the-shelf clients, such as laptops or smartphones.

The rest of the paper is organized as follows. Section II describes related work. Section III presents the proposed network architecture. Section IV describes our approach to solve the optimization problem. Sections V and VI describe our implementation of SmartAP and its experimental evaluation. The conclusions of the paper are drawn in Section VII.

II. Related Work

The first work that proposed the use of virtualization as a means for enabling simultaneous connections to multiple WLAN networks was Multinet [5]. Juggler [6] provides a description of an implementation that reduces the switching time to 3 ms to enable fast switching. However, the closest systems to SmartAP are FatVAP [2] and THEMIS [3], which however perform WLAN backhaul bandwidth aggregation using single-radio clients. FatVAP optimizes the system to achieve maximum throughput per client, but leads to unfair bandwidth distribution in certain scenarios that could refrain users from participating into the sharing scheme. THEMIS solves this problem by guaranteeing a fair sharing of the backhaul bandwidth and it offers the possibility of providing price-based priorities to enable commercial deployments. The main shortcoming of [2] and [3] is that they require client-side modifications, which makes their solution unpractical for commercial deployment. SmartAP addresses this key issue and it provides a way to achieve backhaul bandwidth aggregation modifying only a subset of devices (the APs) that can be controlled by the network operator.

The use of WLANs to share backhaul resources has been discussed extensively [7]–[10]. In [7], the authors propose to normally utilize the wire line connection to the home AP and increase the uplink capacity by connecting to neighboring APs and aggregating their backhaul capacity. Wireless Distribution System (WDS), which enables inter-AP communication could be used to perform backhaul bandwidth aggregation of APs in the same channel. However, using APs with the same radio-channel does not scale as the throughput obtained by each client in a wireless network decreases with the number of nodes [11]. [10] is instead a solution where both APs and clients must be modified to exchange signaling information.

In addition, commercial initiatives like FON [4] and Meraki [12] aim to create nation-wide WiFi sharing communities that offer the possibility for users to connect to the Internet whenever they are in range of an access point (AP) from the community. However, none of these approaches has addressed the potential gain of sharing any unused AP backhaul resource in a neighborhood for bandwidth aggregation.

III. SmartAP Architecture

Our goal is to make WLAN backhaul aggregation practically deployable. To this end, we require a solution that must fulfill the following four requirements:

- **AP-based solution:** a commercially viable solution requires modifications to the minimum number of devices possible.
- **Single-radio APs:** the solution must be cost-effective without extra hardware needed.

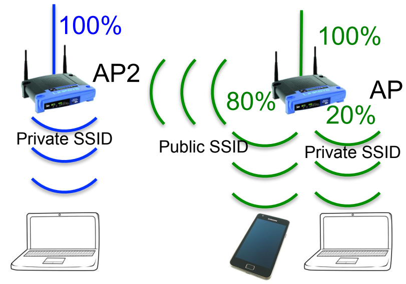

Fig. 1. Example of an SmartAP network.

- **Unmodified WiFi clients:** in order to provide aggregation to all types of WiFi devices, we require no modifications to off-the-shelf WiFi clients.
- **Backhaul connections primarily dedicated to serve owner's devices:** only aggregate unused bandwidth from neighboring APs.

The high level architecture, depicted in Fig. 1, is that of a community that enables its subscribers to connect to any AP participating in the sharing scheme. Differently from traditional WiFi sharing, such as FON [4], in our architecture there are not only users that connect to participating APs, but also other APs in order to make use of the additional backhaul capacity. As such, SmartAP is designed to switch between two modes:

- Sharing mode (SM): represented by AP1 in Fig. 1. In this mode, SmartAP offers the bandwidth that its owners are not using to all the participants in the sharing scheme.
- Aggregation mode (AM): represented by AP2 in Fig. 1. In this mode SmartAP does not have unused backhaul capacity and looks for neighboring SmartAPS to aggregate their backhaul.

The example in Fig. 1 shows two SmartAPs, each one in a different state. The owner of AP1 is using 20% of its backhaul capacity. Thus, AP1 is sharing the remaining 80% with other participants in the sharing scheme. In contrast, the owner of AP2 is saturating its backhaul and requires extra capacity if available. In this situation, AP2 connects to AP1 and redirects to its clients the backhaul capacity that the owner of AP1 is not using. This 80% of capacity that AP1 shares is split between the smartphone that is away from home and the clients of AP2.

SmartAP's mode of operation depends on the utilization of its wired backhaul link. A low wired backhaul link utilization implies that SmartAP operates in SM, and shares its available wired backhaul bandwidth with nearby APs. Conversely, a high AP backhaul utilization causes SmartAP to operate in AM, thereby aggregating the bandwidth of nearby APs in SM. Next, we describe SM and AM in more detail.

a) Sharing mode: In SM, the unused backhaul capacity is shared by means of a public SSID to which all the participants in the sharing scheme can connect (e.g. neighboring SmartAPs and users that are away from their AP). The owners of the AP connect to a private SSID and their traffic is separated from the shared traffic. Since the wired backhaul capacity

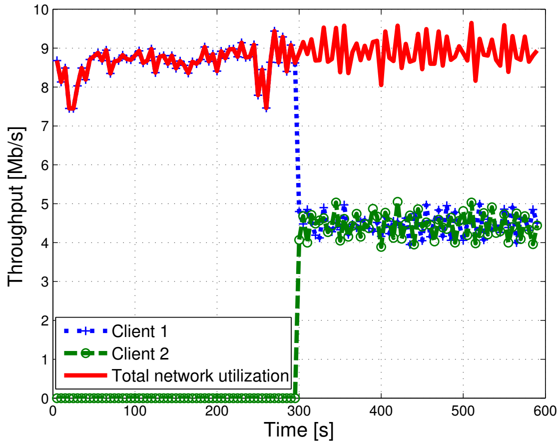

Fig. 2. Backhaul aggregation with three APs and two clients.

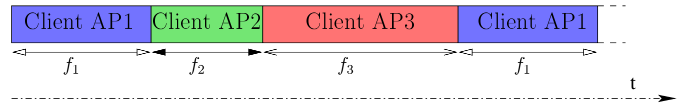

(a) Wireless client-based TDMA aggregation cycle.

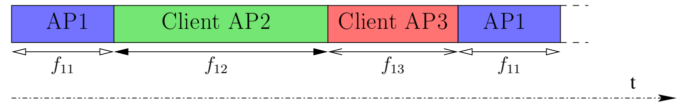

(b) Wireless AP-based TDMA aggregation cycle (SmartAP).

Fig. 3. Single-radio multi-channel virtualization.

is not virtualized, we prioritize the private SSID over the public SSID. SmartAP broadcasts its available-for-aggregation capacity so that any neighboring SmartAP in AM mode can decide whether to use it or not. This value is

$$b - u_{priv} - \gamma \times b,$$

where b is the AP's backhaul capacity, u_{priv} is the utilization by the clients of the private SSID and $\gamma \in (0, 1)$ is a parameter used to prevent SmartAPs from connecting to APs that are close to switching to Aggregation mode. SmartAPs use beacon frames to broadcast this value.

b) Aggregation mode: In this mode, the owners of the SmartAP require its full backhaul capacity. Therefore, SmartAP does not share its backhaul bandwidth (no public SSID broadcasted). In order to improve the throughput provided to its clients, it scans the wireless medium to see the status of neighboring APs (e.g. if they broadcast the public SSID) and becomes a client of the ones that are available[1].

A. An Example

To portrait the key features of SmartAP's architecture, we run an experiment in a three household scenario. In this particular case, there are three SmartAPs within wireless range of each other. Initially, the network features a single client that connects to its home AP (AP1). Given that the traffic requirements exceed the capacity of the backhaul link, AP1 connects to AP2 and AP3 to access more WLAN backhaul bandwidth. At time 300 seconds, a second client appears in the network, and connects to its home AP (AP3). Given that both clients require more throughput than their backhaul links can provide, AP3 switches to AM mode, and AP1 and AP3 use AP2 for assistance. Fig. 2 presents the throughput of the clients connected to AP1 and AP3 as a function of time in the aforementioned topology. At the beginning, the client of AP1 is able to aggregate the capacity of all three backhaul links of 3 Mb/s each. When the second client arrives, the system is able to provide a fair share of the backhaul capacity of AP2, leading to an effective throughput of 4.5 Mb/s for each client.

To realize the full potential of backhaul bandwidth aggregation, SmartAPs must be able to act as clients of APs operating in different radio channels. Therefore, as one of our requirements is to have single-radio APs, SmartAP must connect to APs in different radio frequencies, as we show next.

B. Single-radio multi-channel operations at the AP

Bandwidth aggregation in the state of the art client-based schemes [2], [3], [14] is done through Time Division Multiple Access (TDMA). As shown in Fig. 3(a), the wireless client sequentially connects to all selected APs within range in a round robin fashion over cycles of 100 ms. A client signals to the APs that it is going to enter Power Saving Mode (PSM) before switching to the next AP in the cycle. The percentage of time devoted to each AP (f_i) is determined by the capacity of the wireless channels and their backhauls in order to maximize a utility function.

Fig. 3(b) shows the TDMA cycle of AP-based aggregation systems. While client-based systems need to optimize a cycle of station mode operations, AP-based systems such as SmartAP have to account for AP and station mode operations on the same hardware. SmartAP must spend the appropriate amount of time (f_{ij}) on its AP radio-frequency to forward the data it has collected both over its wired backhaul and the neighboring APs backhauls to its clients. In the next section we present our solution to address this problem.

IV. Formulation

In this section, we describe our approach to solve the problem of selecting the percentage of time that SmartAP connects to either its clients (f_{11}) or APs in range (f_{1j}). Our approach is to map the AP-based aggregation scheme to the client-based scheme. This allows us to apply any available optimization algorithm, such as [2], [3].

The optimization of client-based schemes requires that all clients estimate the following parameters: i) wireless capacity[2] (ω_i), ii) AP backhaul capacity (b_i) and iii) AP backhaul utilization (u_i). Fig. 4(a) and Fig. 4(b) show the network topologies

[1] [3], [13] show that, in typical residential environments, clients are in range of $4-5$ APs with 802.11 links that offer more than 5 Mb/s each.

[2] The wireless capacity refers to the maximum throughput that can be obtained using that link when using it 100% of the time and takes into account the existing interference in the wireless link.

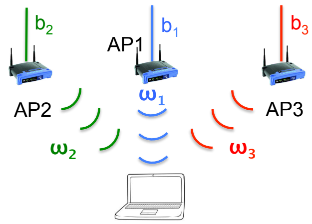

(a) Client-based Aggregation.

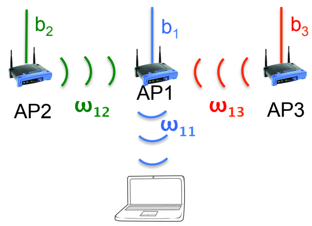

(b) AP-based Aggregation.

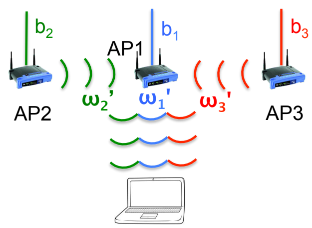

(c) Overlay wireless channels.

Fig. 4. Network topologies in a scenario with three APs and one client.

of both client- and an AP-based aggregation systems. We observe that the two schemes have different wireless links. The main difference is that, in our AP-based scheme, data coming from the backhaul of AP2 must traverse two wireless links with capacities ω_{12} and ω_{11}. The end-to-end wireless capacity resulting from these two hops, ω'_2, is represented by the *overlay wireless channel* depicted in Fig. 4(c).

While the mapping between ω_{11} and ω'_1 is direct, $\omega_{12} \neq \omega'_2$ and $\omega_{13} \neq \omega'_3$. Let us focus on ω'_2, which can be obtained using the time required to transmit a frame of size P bits through both hops:

$$\text{tx\_time} = \frac{P}{\omega_{12}} + \frac{P}{\omega_{11}} = P\frac{\omega_{11} + \omega_{12}}{\omega_{11}\omega_{12}}, \tag{1}$$

Then,

$$\omega'_2 = \frac{\omega_{11}\omega_{12}}{\omega_{11} + \omega_{12}}. \tag{2}$$

The value of ω'_3 can be obtained likewise. Now we have the same ingredients (ω'_i, b_i and u_i) used in [3] to obtain the percentage of time to allocate to each backhaul: f'_i. Then, the result of applying the client-based formulation is the percentage of time that AP1 needs to spend on each overlay wireless link: ω'_i. However, we need the percentage of time that AP1 must devote to the real wireless links: ω_{ij}.

Note that f'_1 accounts just for the time AP1 needs to devote to transmit b_1 Mb/s to its client and does not consider the time to transmit data collected over neighboring APs. Let us focus on the two-hop link ω'_2. The value of f'_2 is the percentage of time to devote to the overlay wireless channel that connects the client of AP1 with AP2. Then, the mapping between f'_2 and f_{12} can be computed from the percentage of time that a packet of size P spends on each of the two links, given ω_{11} and ω_{12}.

$$(\text{\% of time on link } 1-2) = \frac{\omega_{11}}{\omega_{11} + \omega_{12}} \tag{3}$$

Then, the percentage of time connected to AP2 is:

$$f_{12} = f'_2\frac{\omega_{11}}{\omega_{11} + \omega_{12}}. \tag{4}$$

while the percentage of time spent as AP is:

$$f_{11} = f'_1 + f'_2\frac{\omega_{12}}{\omega_{11} + \omega_{12}} + f'_3\frac{\omega_{13}}{\omega_{11} + \omega_{13}} \tag{5}$$

We can then generalize the above equations for smartAP i and $N-1$ neighboring APs as

$$f_{ij} = \begin{cases} f'_j \dfrac{\omega_{ii}}{\omega_{ii} + \omega_{ij}} & j \neq i \\ f'_i + \displaystyle\sum_{k=1,k\neq i}^{N} f'_k \dfrac{\omega_{ik}}{\omega_{ii} + \omega_{ik}} & j = i \end{cases} \tag{6}$$

Table I shows the values of f'_i and ω'_i for the example shown in Fig. 4(b) and three APs, and considering three backhaul links with capacity $b_i = 1$ Mb/s.

V. Implementation

In this section we describe the implementation of SmartAP.

A. Implementation

We have implemented SmartAP in a desktop computer with linux kernel 2.6.32 and an Atheros-based 802.11 PCI Express card controlled by the ath9k [15] Linux driver. The ath9k code of compat-wireless-2.6.32 has an initial implementation of multi-channel virtualization. We modified that code to enable TDMA scheduling, while limiting any associated loss in performance. Our main changes are the following:

- We implement a silencing mechanism by which SmartAPs can silence the contention domain to mask the disappearance from their channel of operation to their clients (Section V-B).
- In order to start the switching process, we stop the upper layer queues, drain the hardware queue and control that the last frame sent before switching is the one indicating that SmartAP is going to PSM.
- We reduce the hardware queue length to increase the granularity of the switching events.

TABLE I
Example of the mapping between the three "overlay channels" and the real wireless channels.

	AP_1	AP_2	AP_3
b_i [Mb/s]	1	1	1
ω_{ij} [Mb/s]	$\omega_{11} = 20$	$\omega_{12} = 10$	$\omega_{13} = 10$
ω'_i [Mb/s]	$\omega'_1 = 20$	$\omega'_2 = 6.67$	$\omega'_3 = 6.67$
f'_i	$f'_1 = 0.05$	$f'_2 = 0.15$	$f'_3 = 0.15$
f_{ij}	$f_{11} = 0.15$	$f_{12} = 0.10$	$f_{13} = 0.10$

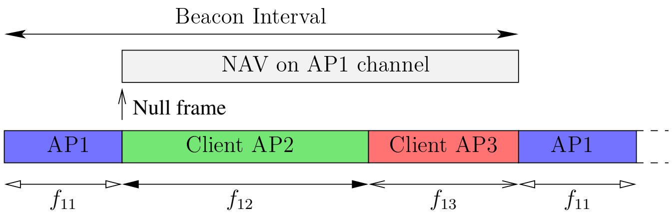

Fig. 5. SmartAP sends a null frame to silence its clients for a percentage of time $\sum_{j=1,j\neq i}^{N} f_{ij}$.

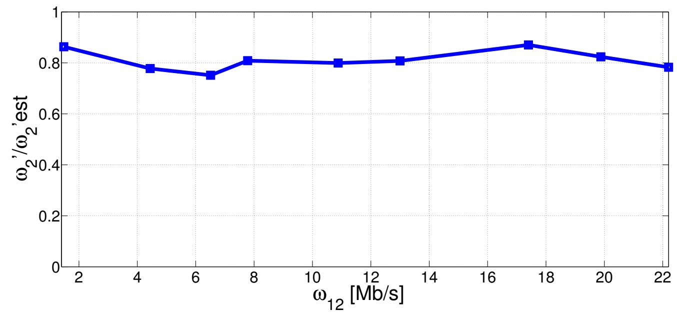

Fig. 6. Accuracy of ω_2' estimator.

B. Methodology to leave the AP channel

The main challenge of the implementation is to avoid any losses caused by the AP leaving its channel to aggregate the bandwidth from other APs in range. In percentage of time, SmartAP i must leave its AP channel for $\sum_{j=1,j\neq i}^{N} f_{ij}$. Differently from client-based approaches, APs cannot enter PSM, and they are supposed to be always active to collect the 802.11 frames from their clients.

In order to resolve this issue, we use standard features of 802.11 for compatibility and proceed as follows. Prior to transmitting a frame, any 802.11 device estimates the time it will take to send it, given the transmission rate and frame length, and writes this value into the duration field of the MAC header. All devices in range update their NAV (network allocation vector) according to the duration field of the frames they receive. As a result, none of the 802.11 devices in range will attempt to access the medium until the NAV expires. SmartAP uses the NAV to reserve the channel for the amount of time it will be out of its AP channel.

We refer to Fig 5. SmartAP sends a null frame before leaving its primary channel of operation with a dummy MAC address as receiver[3], and with a duration field equal to the amount of time it will be acting as a client for all the other APs. The maximum value of the duration field is ≈ 32 ms (15 bits to indicate the duration in microseconds). This process can be iterated: One more dummy frame can be sent to the clients at the expiration of this silencing time to extend the time required by SmartAP to collect bandwidth from other APs. SmartAP completes this process in one Beacon interval, and it is ready to send the next Beacon frames as expected by its clients. As desired in our design requirements of Section III, our clients remain totally unmodified, thus allowing to easily deploy the solution and perform experiments with commodity laptops, as well as smartphones.

C. Algorithm and Estimators for AP-based Optimization

We use the algorithm presented in [3] and summarized in the Appendix to select the percentage of time f_i', and the formula given in Eq. 4 and Eq. 5 for the percentage of time f_{ij}. Since our formulation allows to directly map AP-based into client-based approaches, other algorithms could be also used. Note also that these formulas also take into account the switching cost. The values required to perform the optimization are: ω_{ii}, ω_{ij}, b_i and u_i. As explained in Section III, the available backhaul capacity of neighboring APs (b_i) and their backhaul utilization (u_i) is announced through beacon frames. Finally, the wireless capacities are obtained using the estimators presented in [3].

[3]If the channel reservation frame is instead transmitted to one of the clients of the SmartAP, this client might start transmitting at the end of the reception of the frame.

VI. EXPERIMENTAL EVALUATION

In this section we evaluate the performance of SmartAP on a small scale testbed. Our evaluation focuses on: i) the assessment of two-hop wireless capacity estimation described in Section IV, ii) the performance gain of SmartAP in a realistic setup, and iii) an analytical comparison of the aggregation capacity of AP- versus client-based aggregation schemes.

A. Experimental Setup

Our testbed consists of two SmartAPs, one Linksys WRT54GL AP and three machines acting as servers on the wired network. To study the impact of different backhaul capacities on the performance of SmartAP, we use traffic control (tc) to rate limit the backhaul links. To study the impact of different wireless capacities, we modify the transmission rate of APs. We use Dell laptops (latitude D620) running ubuntu 8.04 (Linux kernel 2.6.24) to connect to the SmartAPs to perform the majority of the throughput experiments we show in this paper. In addition, we have tested Android phones (e.g. we have tested an HTC Nexus One and a Samsung Galaxy SII) connecting to the SmartAPs and obtained the same performance as that obtained by the laptops. We establish downlink TCP connections with Iperf[4] to measure the throughput achieved in each scenario. All experiments were conducted during night to avoid interference from networks that do not belong to the test bed.

B. Two-hop capacity validation

In Section IV we showed that it is possible to map the formulation of the AP-based aggregation problem to the client-based one and, hence, apply any solution from the ones already proposed in the literature. The formulation requires the two-hop wireless capacity, which is assessed in this section.

We then study the accuracy of ω_2' in eq. 2. For the evaluation, we measure ω_{12} for all the transmission rates

[4]http://iperf.fr/

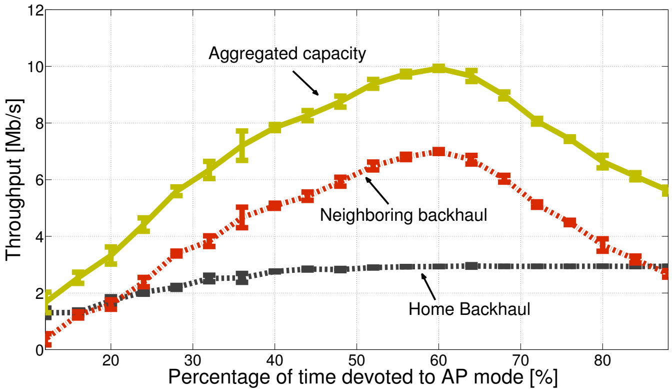

Fig. 7. Aggregation capacity with a home backhaul of 3 Mb/s.

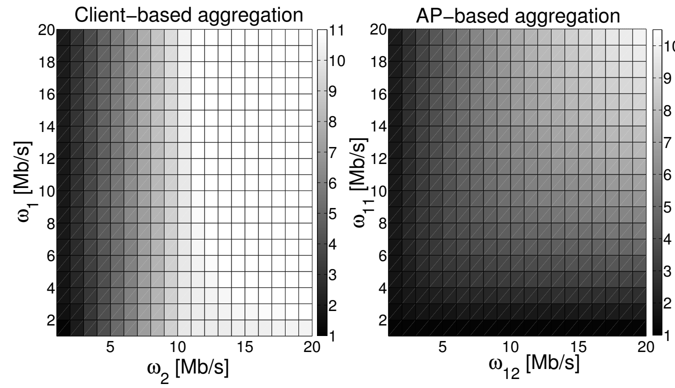

Fig. 8. Aggregation capacity of client- and AP-based systems with a home backhaul of 1 Mb/s and a neighboring backhaul of 10 Mb/s.

of an 802.11g AP[5]. This emulates the effect of having the neighboring AP farther away with lower wireless capacity. In order to measure ω_{12}, we performe five iperf tests of 50 seconds each between the neighboring AP and SmartAP, using only one virtual interface in client mode. We also measure the relaying capacity of SmartAP, varying the percentage of time devoted to the AP interface. The maximum relaying capacity in each of these tests corresponds to the experimental ω'_2. In addition, we measure ω_{11} = 22.2 Mb/s and compute the estimated ω'_2 for each of the experiments using eq. 2.

Fig. 6 shows the result of dividing the experimental ω'_2 over the estimated ω'_2 (indicated as ω'_2est in the figure). We observe that the experimental values of ω'_2 are 0.8 times the estimated values, which corresponds to the impact of fast channel switching on the throughput resulting from our implementation. Taking into account the cost of switching, the values of ω'_2 are the ones we compute using the change of variables. This validates the mapping between SmartAP and client-based aggregation schemes.

C. Aggregation capacity

To explore the maximum aggregation capacity of SmartAP, we use a network with two APs and one client. We set the capacity of the home backhaul (b_1) to 3 Mb/s and impose no limitation to the neighboring backhaul (b_2) or the wireless channel between SmartAP and the neighboring AP — AP2 transmits at a physical rate of 54 Mb/s —. In this scenario SmartAP provides a total throughput of 9.9 Mb/s to its client: 3 Mb/s coming from the home backhaul and 6.9 Mb/s gathered from the neighboring AP. The optimal operating point of SmartAP is to devote 40% of the time to collect bandwidth from the neighboring AP and the remaining 60% to serve the aggregated throughput to its client.

Fig. 7 shows the evolution of the total throughput obtained by the client depending on the percentage of time that SmartAP acts as an AP. When the time devoted to the AP interface is lower than the optimum value, the throughput obtained from the home backhaul is lower than 3 Mb/s. This is because the bottleneck in this situation is the link between SmartAP and its client and all the TCP flows will compete to obtain the available capacity. Instead, when SmartAP allocates more than 60% of the time to its AP interface, it has not enough time acting as a client to collect the bandwidth from the neighboring AP backhaul.

D. Client- vs AP-based bandwidth aggregation

Previous sections have shown the feasibility of AP-based aggregation systems and their benefits. To analyze the potential of SmartAP with respect to client-based aggregation systems, we perform an analytical comparison of its aggregation capacity versus client-based systems in a simple yet effective scenario: one client in proximity to its AP and both of them in range of a neighboring AP with similar quality ($\omega_2 = \omega_{12}$). In the test, the backhaul capacity of the home AP (AP1) is 1 Mb/s while the neighboring AP (AP2) has a 10 Mb/s backhaul capacity.

Fig. 8 shows the throughput that each aggregation scheme provides to the client. Comparing the results of both systems, we observe that the client-based solution is able to aggregate the full capacity of the network (11 Mb/s) while the AP-based scheme achieves a maximum of 10.5 Mb/s. As the backhaul capacity of AP2 is much higher than that of AP1, client-based solutions can obtain high throughput when they can connect to AP2 with good quality, and regardless of the wireless capacity of AP1. In contrast, AP-based schemes require that both links have high wireless capacity to obtain all the backhaul bandwidth of the network. Nevertheless, respect to absence of aggregation, SmartAP offers four times higher throughput to its clients when $\omega_{12} > 5$ Mb/s and $\omega_{11} \sim 20$ Mb/s.

In this scenario, the home AP1 backhaul contributes 9% of the total capacity. This implies that the benefit of obtaining this capacity is minimal compared to the neighboring capacity. The result is that when the link to the neighboring AP offers higher capacity than the home backhaul, the maximum throughput is obtained by connecting to the neighboring AP (as long as this bandwidth is not used by the owner). Further reducing the capacity of the home backhaul would lead to a situation

[5]We do not show the results for the transmission rates of 11 Mb/s and 5.5 Mb/s as they are similar to the ones obtained for 12 Mb/s and 6 Mb/s.

in which SmartAP becomes a relay of the neighboring AP. In case the home AP is in between the client and the neighboring AP, the client will obtain higher throughput using the home AP as a relay than trying to directly connect to the neighboring AP.

VII. Conclusion

In this paper, we have addressed a fundamental problem with prior WLAN bandwidth aggregation solutions: their need for client modifications makes their deployment cost prohibitive. Our solution is SmartAP. This new approach requires only software modifications at the APs, controlled by the network operator, and it does not require either any software or hardware change in the WLAN clients. We have studied the problem formulation and shown that the approach can be mapped to the original problem statement for modified clients. We have then implemented a prototype of our approach using a customized open-source driver for Atheros chipsets. Our extensive experimental tests have shown that SmartAP obtains substantial throughput gain, while being less efficient in terms of aggregation respect to the state-of-the-art client-based approach. Tests have been conducted using a variety of unmodified WiFi devices, such as laptops and different Android phones. SmartAP is ready to be commercially deployed with simple remote firmware upgrades of off-the-shelf APs.

Acknowledgments

This research was funded by the European Communitys Seventh Framework Programme (FP7/2007-2013) no. 258378 (FIGARO). The authors would like to thank Konstantina Papagiannaki for her valuable feedback.

References

[1] Ministry of Information and Communication Technologies of Colombia, "Report of Internet statistics in Colombia, 2011." Online available at http://goo.gl/t33ft, viewed 23-09-2013.

[2] S. Kandula, K. C.-J. Lin, T. Badirkhanli, and D. Katabi, "Fatvap: aggregating ap backhaul capacity to maximize throughput," in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'08, (Berkeley, CA, USA), pp. 89–104, USENIX Association, 2008.

[3] D. Giustiniano, E. Goma, A. Lopez Toledo, I. Dangerfield, J. Morillo, and P. Rodriguez, "Fair wlan backhaul aggregation," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, (New York, NY, USA), pp. 269–280, ACM, 2010.

[4] "FON." http://www.fon.com/, viewed 23-09-2013.

[5] R. Chandra, P. Bahl, and P. Bahl, "Multinet: connecting to multiple ieee 802.11 networks using a single wireless card," in *INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies*, vol. 2, pp. 882–893 vol.2, 2004.

[6] A. Nicholson, S. Wolchok, and B. Noble, "Juggler: Virtual networks for fun and profit," *Mobile Computing, IEEE Transactions on*, vol. 9, no. 1, pp. 31–43, 2010.

[7] S. Jakubczak, D. G. Andersen, M. Kaminsky, D. Papagiannaki, and S. Seshan, "Link-alike: using wireless to share network resources in a neighborhood," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 12, no. 4, pp. 1–14, 2008.

[8] X. Ai, V. Srinivasan, and C.-K. Tham, "Wi-Sh: A Simple, Robust Credit Based Wi-Fi Community Network," in *Proc. of the IEEE INFOCOM Conf.*, April 2009.

[9] D. Giustiniano, E. Goma, A. Lopez Toledo, and G. Athanasiou, "Optimizing TCP performance in multi-AP residential broadband connections via minislot access," *Journal of Computer Networks and Communications*, vol. 2013, 2013.

[10] E. Chai and K. G. Shin, "Sidekick: Ap aggregation over partially overlapping channels," in *Network Protocols (ICNP), 2011 19th IEEE International Conference on*, pp. 301–310, IEEE, 2011.

[11] P. Gupta and P. Kumar, "The capacity of wireless networks," *Information Theory, IEEE Transactions on*, vol. 46, pp. 388 –404, mar 2000.

[12] "Meraki." http://www.meraki.com/, viewed 23-09-2013.

[13] N. Thompson, G. He, and H. Luo, "Flow Scheduling for End-host Multihoming," in *Proc. of INFOCOM'06*, 2006.

[14] D. Giustiniano, E. Goma, A. Lopez, and P. Rodriguez, "WiSwitcher: an efficient client for managing multiple APs," in *Proceedings of the 2nd ACM SIGCOMM workshop on Programmable routers for extensible services of tomorrow*, pp. 43–48, ACM, 2009.

[15] Linux Wireless, "ath9k linux driver." http://goo.gl/pg89Y, viewed 23-09-2013.

Appendix

Let $\mathcal{S}$ be the set of stations and $\mathcal{A}$ the set of APs. Denote T_{ik} as the throughput sent from AP_i to station k. And let $y_k = \sum_{i\in\mathcal{A}} T_{ik}$ denote the total throughput received by station k. Let $U(\cdot)$ be a differentiable, strictly concave, increasing function which represents the utility at every station as a function of the received throughput. The fairness problem is modeled as

$$\max \sum_{k\in\mathcal{S}} U(y_k) \tag{7}$$

$$\text{s. t.} \quad \sum_{k\in\mathcal{S}} T_{ik} \le b_i, \quad \forall i \in \mathcal{A}, \tag{8}$$

$$\sum_{i\in\mathcal{A}, w_{ik}>0} \frac{T_{ik}}{w_{ik}} \le 1, \forall k \in \mathcal{S}, \tag{9}$$

$$T_{ik} \ge 0, \quad \forall i \in \mathcal{A}, \forall k \in \mathcal{S}, \tag{10}$$

where w_{ik} is the wireless capacity at which station k can receive from AP_i ($w_{ik} = 0$ if station k does not connect to AP_i), that takes into account the interference from other clients connected to that AP, and b_i is the backhaul capacity of AP_i.

In [3], authors use a primal-dual formulation using a gradient descent algorithm to solve the above optimization deriving the following optimal rate update rule

$$T_{ik} = \hat{T_{ik}} + \alpha \left(U'(y_k) - p_i - q_{ik} \right), \tag{11}$$

where $\hat{T_{ik}}$ is the bandwidth request in the previous step of the algorithm, $U'(y_k)$ is the derivative of the utility function evaluated at the current throughput received by the station y_k, and α is the step size of the rate update algorithm. The quantities p_i and q_{ik} are the prices corresponding to constraints (8) and (9) respectively.

The percentage of time that each client must allocate to each AP results from enforcing the obtained bandwidth requests using the following equation:

$$f_{ik} = \sigma_{ik} \frac{T_{ik}}{w_{ik}} + c_i, \tag{12}$$

where σ_{ik} is the correction factor, and c_i is the overhead of switching from one AP to the next.