

Fair WLAN Backhaul Aggregation

Domenico Giustiniano^{†*}, Eduard Goma[†], Alberto Lopez Toledo^{†||}, Ian Dangerfield^{‡¶},
Julian Morillo^{§¶}, and Pablo Rodriguez[†]

[†]Telefonica Research, [‡]Hamilton Institute, [§]Universitat Politècnica de Catalunya
{domenico,goma,alopez,t,pablorr}@tid.es, ian.dangerfield@nuim.ie,
jmorillo@ac.upc.edu

ABSTRACT

Aggregating multiple 802.11 Access Point (AP) backhauls using a single-radio WLAN card has been considered as a way of bypassing the backhaul capacity limit. However, current AP aggregation solutions greedily maximize the individual station throughput without taking fairness into account. This can lead to grossly unfair throughput distributions, which can discourage user participation and severely limit commercial deployability.

Motivated by this problem, we present THEMIS, a single-radio station that performs multi-AP backhaul aggregation in a fair and distributed way, without requiring any change in the network. We implement THEMIS on commodity hardware, evaluate it extensively through controlled experimental tests, and validate it in a deployment spanning 3 floors of a multistory building. THEMIS is being used in a commercial trial by a major broadband provider to its customers.

Categories and Subject Descriptors

C.2.5 [Computer Communication Networks]: Local and Wide-Area Networks—*Access schemes*; C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Design, Experimentation, Performance

1. INTRODUCTION

In urban environments, residential users can potentially see multiple 802.11 APs in range with high quality [1], usually connected to broadband links. As the speeds of 802.11

[¶]This work was carried out while the authors were at Telefonica Research Barcelona, Spain.

*Now at Disney Research: domenico@disneyresearch.com.

^{||}Supported by the Institutió Catalana de Recerca i Estudis Avançats (ICREA).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'10, September 20–24, 2010, Chicago, Illinois, USA.

Copyright 2010 ACM 978-1-4503-0181-7/10/09 ...\$10.00.

WLAN are typically an order of magnitude higher than those of standard broadband connections, one can use a single 802.11 wireless card to aggregate the bandwidth of multiple AP backhauls in range by virtualizing the card and cycling over the APs in a TDMA fashion. The result of such multi-AP aggregation scheme is that stations will connect to several APs in range and share their backhaul connections.

In that scenario, using an aggregation scheme like Fat-VAP [2], where stations greedily maximize their individual throughput, may lead to severe unfair situations. Fairness is important because it can impact individual users performance and reduce its applicability on a commercial setting. For example, a station that is unluckily located in an area with only one AP in range, can see its throughput significantly lowered by other stations sharing the same AP, even if those stations could get spare bandwidth from other APs. This is what we call *topology unfairness*. We argue that providing a fair distribution of throughput even in such heterogeneous situations is crucial to maintain a certain level of service across all users. Without some form of fairness, the perceived value of the system is severely reduced, and users will not participate.

Other unfairness situations also exist. For instance, stations using applications with many TCP flows, such as P2P, can severely affect the performance of other stations running single-flow applications such as Web downloads. We call this situation *flow distribution unfairness*, and can result in some stations obtaining much less throughput than what they would obtain without sharing.

Another example of unfairness could appear in a scenario where customers with different subscription plans share their broadband links. For instance, fast broadband customers (that pay more than slow broadband customers), should obtain a greater share of the spare backhaul capacity. If this is not enforced, customers may be inclined to buy slower (and cheaper) broadband connections and free-ride on their neighbors' spare bandwidth. This is a typical "tragedy of the commons" example: people tend to over-exploit the shared resource by minimizing their contribution (their broadband contracted speed), ultimately cannibalizing the shared resource. Moreover, this eliminates the incentive of an ISP to deploy the sharing system, because it threatens its business model. We call this *billing unfairness* (Section 2).

The above fairness scenarios can have a dramatic impact on the deployability of various multi-AP aggregation schemes including: a) community-based sharing schemes (e.g. FON [3], Wi-Sh [4]), b) Telco-managed sharing schemes where residential WiFi gateways are shared across all users

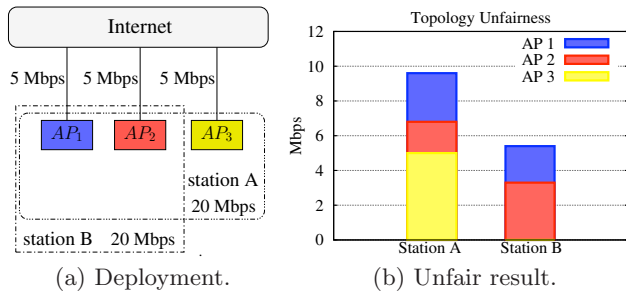


Figure 1: Unfairness for users with different AP connectivity.

that subscribe to the service, and c) commercial AP aggregation scenarios (e.g. airport hotspots). Moreover, existing aggregation schemes such as FatVAP [2] and VirtualWiFi [5] are not designed with fairness in mind, and hence cannot be directly applied to the above scenarios.

Motivated by these problems, we introduce THEMIS¹, a single-radio station that fairly aggregates the backhaul bandwidths of several APs. We extensively evaluate THEMIS in controlled scenarios, and show that it provides a fair distribution of the available backhaul bandwidth among users (Section 4). Finally, we validated THEMIS by emulating a typical urban neighborhood environment consisting of a setting of 10 commercial ADSLs with their correspondent 802.11g APs over 3 consecutive floors of a multistory building (Section 5).

2. FAIR WIRELESS BACKHAUL AGGREGATION

Let us consider the multi-AP backhaul aggregation system depicted in Fig. 3, where single-radio 802.11 stations simultaneously connect to one or more APs. In this scenario, the AP backhaul bandwidth of the APs is shared among the stations. Next we will give some illustrative examples which show the need for fairness and how greedy schemes, such as [2], fail.

Topology unfairness. Consider the experiment² depicted in Fig. 1(a), where stations A and B share 3 APs, each of them having a 5 Mbps backhaul. The wireless speed from each station to the three APs is 20 Mbps. However, because of its location, station B has only two APs in range, while station A can reliably connect to all the APs. A fair distribution of the aggregated AP backhaul would assign half of the backhaul capacity — 7.5 Mbps — to each station. Using a throughput maximization scheme as in [2], station B obtains 5 Mbps, almost half of the throughput of station A, which obtains more than 9 Mbps due to its better location (Fig. 1(b)).

Flow distribution unfairness. Consider now the experiment in Fig. 2(a), where stations A and B connect to two APs with 5 Mbps backhauls. The wireless speed between the stations and the APs is 20 Mbps. Station B starts one down-

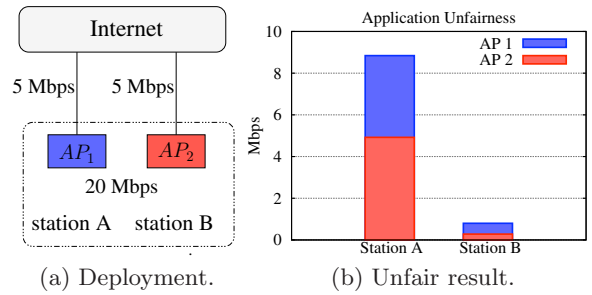


Figure 2: Unfairness for users with different number of flows.

load per AP, each using a single TCP flow. Station A, on the other hand, starts one download per AP, but each using 10 TCP flows. The experiment is set up to guarantee that the flows are not limited by the end-to-end connection, i.e. the bottleneck is in the AP backhaul. In this scenario, a fair distribution of the AP backhaul would result in each station receiving 5 Mbps. However, if the stations aim to maximize their individual aggregate throughput without taking fairness into account as in [2], the result is a gross unfair distribution of the bandwidth, with station A receiving almost 9 Mbps, most of the aggregated bandwidth, while station B receives less than 1 Mbps (Fig. 2(b)). A similar scenario could be shown for the case of billing unfairness.

The above examples clearly illustrate the need to provide a fairness mechanism for the multi-AP backhaul aggregation scheme. However, it is important to agree on some notion of fairness, since each one could have different design implications and trade-offs. We discuss this in detail next.

2.1 What Kind of Fairness?

In order to address the unfairness situations described above, we start by describing our fairness requirements. First we would like to ensure that fairness is achieved at the level of the station’s total received throughput, as opposed to individual flows or packet level fairness (**per-station fairness**). Second, we would like to ensure that users with better subscription plans (e.g. faster broadband links) obtain greater share of the aggregated AP backhaul bandwidth than users with cheaper subscription plans. Thus, in the examples above the throughput should be obtained proportionally to their priority (**weighted fairness**). Third, fairness should be enforced across all shared APs, and not just at the single AP level to ensure a fair global throughput allocation (**across-AP fairness**). Fourth, we want to provide a fairness scheme that is efficient in terms of network utilization and strikes a good balance between fairness and throughput (**efficient fairness**). And finally, we would like to provide a fairness scheme that is stable and has good convergence properties (**stable fairness**). Furthermore, in order to facilitate a wide adoption, we want to minimize the impact on the existing network infrastructure.

There are different reasons why the above requirements cannot be achieved using existing network technologies. For instance, in infrastructure mode, 802.11 does not provide per-station fairness because its downlink behavior is largely dominated by its FIFO packet-level scheduler [6]. TCP, on the other hand, only provides per-flow fairness among competing downlink flows, which is in fact the cause of the flow

¹THEMIS is the Greek goddess of Justice, usually portrayed as an impassive blindfolded woman, holding scales outside a courthouse.

²All the tests and validations in this paper are performed experimentally on realistic scenarios. See Section 4 for details about the experimental setup.

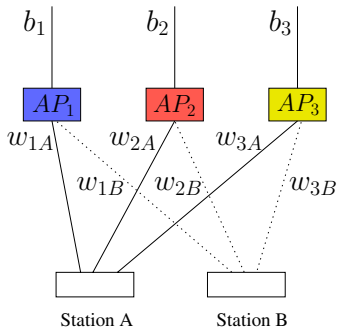


Figure 3: Multi-AP aggregation scenario.

distribution unfairness [7]. Even if one would manage to implement some fairness mechanism at the individual AP level (for example changing the FIFO behavior or introducing some clever time-based scheduler [8]), this would not result in across-AP fairness without the use of explicit signaling among the APs.

Our Choice of Fairness

In wireless systems, it is well known that fairness and throughput are often at odds [9]. For instance, imagine a scenario where two stations are sharing a wireless medium, and their wireless speeds are at a ratio of 10:1. A throughput optimal allocation would only allow the fast station to transmit, because every time slot devoted to the slow station would be wasted in low speed, losing the chance of a fast transmission. At the other extreme, a max-min fair allocation (e.g. one that maximizes the minimum of all station throughputs) would equalize the throughput transmitted by both stations. This allows the slow station to transmit most of the time, causing *performance anomaly* [9], that severely reduces the overall throughput.

Proportional fairness lies in the middle of the two extremes, providing a good compromise between fairness and efficiency (e.g. in [10]). It also achieves a good trade-off in terms of convergence and stability as shown in [11]. Finally, it allows for weighted fairness formulation. *Weighted proportional fairness* meets our efficient, stable and weighted requirements.

To comply with the other two requirements (per station and across-AP fairness), we cannot rely on existing formulations such as those in [2]. In fact [2] uses a knapsack scheduler that maximizes the individual station's throughput, and does not consider how the aggregate throughput is partitioned across stations. As a result, we need a new formulation that takes this problem into consideration. We describe it next.

2.2 Fairness Formulation

Recall the scenario depicted in Fig. 3. 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. We model the fairness problem as³

³For simplicity, and given that current residential traffic is heavily biased towards downloads, our formulation only con-

$$\max \sum_{k \in \mathcal{S}} U(y_k) \quad (1)$$

$$\text{s. t. } \sum_{k \in \mathcal{S}} T_{ik} \leq b_i, \quad \forall i \in \mathcal{A}, \quad (2)$$

$$\sum_{i \in \mathcal{A}, w_{ik} > 0} \frac{T_{ik}}{w_{ik}} \leq 1, \quad \forall k \in \mathcal{S}, \quad (3)$$

$$T_{ik} \geq 0, \quad \forall i \in \mathcal{A}, \forall k \in \mathcal{S}, \quad (4)$$

where w_{ik} is the wireless capacity⁴ at which station k can receive from 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 .

Eq. (2) is the *AP backhaul capacity constraint*, and ensures that the total traffic traversing the AP_i backhaul does not exceed the backhaul capacity b_i . Eq. (3) corresponds to the *station k wireless capacity constraint*, and guarantees that the total traffic received by station k does not exceed the total capacity of its wireless interface. Finally (4) forces the values T_{ik} to be positive.

Note that there exists an additional constraint not included in the formulation, corresponding to the *AP wireless capacity constraint*, namely $\sum_{k \in \mathcal{S}} \frac{T_{ik}}{w_{ik}} \leq 1, \forall i \in \mathcal{A}$. This constraint ensures that the maximum capacity of the wireless interface at AP_i is not exceeded. However, we verified analytically that this constraint may be violated only in the extreme cases of clients severely limited by the wireless. We avoid this situation by preventing stations from connecting to *APs* if their signal-to-noise ratio (SNR) is very low. This makes sense, as a multi-AP aggregation scheme is only useful if the speed of WLAN is greater than the speed of the AP backhaul.

Finally, as described in Section 2.1, we choose a *weighted proportionally fair* utility function $U(y_k) = K_k \cdot \log y_k$, where K_k represents the relative priority of user k (for example, a value linearly dependent to the AP backhaul bandwidth owned by user k). If all the users have the same priority we use $K_k = 1$.

Decomposition and Interpretation

As described in [11], the solution of the above optimization problem can be obtained via the primal-dual formulation using a gradient descent algorithm. From there we derive the following optimal rate update rule

$$T_{ik} = \hat{T}_{ik} + \alpha (U'(y_k) - p_i - q_{ik}), \quad (5)$$

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 algorithm⁵. The quantities p_i and q_{ik} are the prices corresponding to

considers downlink traffic. However an equivalent formulation can be designed for uplink traffic.

⁴Note that $w_{ik} = 0$ if station k does not connect to AP_i .

⁵When using proportional fairness, and in order to reduce oscillations as suggested by [12], we use $\alpha = \alpha' y_k$, with α' the new step size.

constraints (2) and (3) respectively, calculated as follows

$$p_i = \left[\hat{p}_i - \frac{\delta}{b_i} \left(\lambda b_i - \sum_{k \in \mathcal{S}} T_{ik} \right) \right]^+, \quad (6)$$

$$q_{ik} = \left[\hat{q}_{ik} - \frac{\gamma}{w_{ik}} \left(\mu - \sum_{i \in \mathcal{A}} \frac{T_{ik}}{w_{ik}} \right) \right]^+, \quad (7)$$

where \hat{p}_i , \hat{q}_{ik} are the prices obtained in the previous step of the algorithm, and δ and γ are the step sizes of the price update algorithm. In order to improve the network utilization, and as suggested in [12], we normalize the price step size by the link capacities to favor good links. Finally $\lambda, \mu \leq 1$ are the congestion thresholds and $(x)^+ = \max(x, 0)$.

The price p_i in (6) represents the level of congestion on the backhaul of AP_i , and it is a linear function of its available bandwidth. Similarly, q_{ik} in (7) represents the level of congestion on the wireless link from station k to AP_i , and it is a function of the available card time at the station⁶. As congestion increases, the respective prices will increase and the throughput demand T_{ik} of station k through AP_i will decrease according to (5).

The values λ and μ are the *congestion thresholds*, i.e. respectively the level of utilization of the AP_i backhaul and the wireless radio-interface of station k that will trigger the algorithm congestion control. When that happens, the prices p_i and q_{ik} increase, prompting the throughput requests for their respective paths to decrease⁷.

In order to distributedly solve the optimization problem in (1), each station has to periodically obtain the prices (6) and (7) for its links, and then update its rates following (5). However, implementing this algorithm locally at each station without sharing information with the APs and/or other stations has the following challenges:

- once the values T_{ik} in (5) are obtained at station k , those rates need to be enforced at AP_i (Section 3.1).
- in order to calculate the prices p_i in (6) and q_{ik} in (7), each station k needs to obtain the values of b_i and T_{ij} $j \neq k$, which are not directly available at the station. Moreover each station k needs to accurately know the wireless capacity w_{ik} of each AP_i (Section 3.2).
- a single-radio station has to manage the communication with multiple APs on independent radio frequencies. And it has to do it efficiently and using standard-compliant 802.11 (Section 3.3).

Addressing the above challenges in a real system requires careful design and implementation, which we describe next.

3. THEMIS

THEMIS is a single-radio wireless station based on the MadWiFi 0.9.4 driver [13] and the Click modular router 1.6.0 [14], that connects to multiple APs and aggregates their

⁶The time that the card is not being used for transmitting or receiving.

⁷The values of the congestion thresholds represent a performance threshold: the closer to 1 the better if for the network utilization, but the worse is for the short-term fairness of the algorithm.




Figure 4: Time-division access to multiple APs.

backhaul bandwidth. As shown in Fig. 4, THEMIS communicates separately to APs at different radio-frequencies using Time-Division Multiple Access (TDMA). Once connected to one AP, THEMIS transmits and receives traffic according to the 802.11 DCF protocol. The amount of time spent on AP_i is denoted *duty cycle* f_i . The constant time T that THEMIS takes to perform a standard TDMA cycle is called *wireless period*. THEMIS will use any spare duty cycle to do other operations such as AP scanning or saving energy.

3.1 Scheduler

Let us consider a THEMIS station running the optimization algorithm in (1), and calculating the request rate to AP_i to be T_{ik} in (5). In principle, in order to collect the bandwidth T_{ik} from AP_i , station k needs to connect to AP_i during a duty cycle $f_{ik} = T_{ik}/w_{ik}$, where w_{ik} is the wireless capacity from AP_i to station k . By reducing the time spent on AP_i , the duty cycle f_{ik} effectively acts as a gauge that limits the amount of bandwidth that can be received from the AP. As a consequence, TCP flows adjust their transmission rate to meet the request T_{ik} .

There are cases where station k does not receive the expected traffic T_{ik} during the duty cycle f_{ik} . There are various reasons for this discrepancy: wireless losses, congestion in the AP queue, CSMA contention delays in the wireless links, etc. We introduce a correction factor $\sigma_{ik} = T_{ik}/x_{ik}$ to account for the deviation between the expected received traffic T_{ik} and the *actual* traffic x_{ik} received by station k from AP_i during the selected duty cycle f_{ik} . As a result, THEMIS connects to AP_i for

$$f_{ik} = \sigma_{ik} \frac{T_{ik}}{w_{ik}} + c_i, \quad (8)$$

where σ_{ik} is the correction factor, and c_i is the overhead of switching from one AP to the next (see Section 3.3). Note that after applying the correction factor it may happen that the corrected duty cycles exceed the allowed time, violating the station k wireless capacity constraint, i.e., $\sum_{i \in \mathcal{A}} f_{ik} > 1$. In that case we distribute the *wireless period* proportionally among the links as described in the Appendix.

3.2 Estimators

The calculation of the duty cycle f_{ik} in (8) at station k for a given AP_i and the update of the prices p_i and q_{ik} in (6) and (7) require the following information:

- the utilization rate $\beta_i = \sum_{k \in \mathcal{S}} T_{ik}$ of the AP_i backhaul;
- the wireless capacity w_{ik} , that determines the maximum transmission rate of the wireless link; and
- the AP backhaul capacity b_i , that measures the maximum speed at which the AP_i backhaul can send traffic.




Figure 5: THEMIS estimators are based on local processing at MAC and PHY layer.

A straightforward way to obtain those values would be to introduce new signaling to exchange this information between the APs and the stations. However, that would introduce extra overhead, and would also require modifying (or replacing) the existing AP installed base. To avoid it, THEMIS estimates these values locally. Note that it is important to achieve high accuracy on the estimations, because wrong estimations would affect the performance of the scheduler. This is a hard problem because:

- *AP backhaul*: the AP backhaul is shared with other stations and any measure of the AP_i utilization rate β_i and capacity b_i must be done in the limited slice of time $f_{ik} \cdot T$ that station k dedicates to AP_i .
- *wireless link*: the wireless capacity of one AP has to be measured while the AP transmits in saturation. This is not guaranteed because the wireless link is usually not the bottleneck of the end-to-end communication.

We next describe how THEMIS estimates these values.

Utilization Rate of the AP Backhaul

The estimation of the utilization rate β_i of the AP backhaul relies on the fact that every frame sent by an 802.11 AP carries a MAC Sequence Number (SN) in the header. The SN is a module 4095 integer incremented by the AP each time a new frame is sent, and it is independent of the destination. THEMIS stations listen to the traffic sent by AP_i , and store its SNs. By counting the SNs, the THEMIS station knows the amount of packets traversing the AP_i backhaul⁸. Note that this way of counting is robust to packet loss and disconnection periods, as long as the stations do not miss more than 4095 successful frames (retransmitted frames do not increase the SN), which for an average 802.11 frame size would correspond to seconds, an order of magnitude larger than the THEMIS' TDMA period⁹.

Formally, let us refer to Fig. 5. We denote $SN_1^i[\text{First}]$ and $SN_M^i[\text{Last}]$ the MAC sequence number of the first and last packet, respectively, sent by AP_i to any station, during a window of time $M \cdot T$, where M is an integer equal or greater than 1. Then, THEMIS derives the number of packets sent

⁸Here we assume that most of the 802.11 data traffic traversed the AP backhaul, as it is often the case when using 802.11 in infrastructure mode.

⁹To increase the accuracy of the estimation, THEMIS operates in promiscuous mode, thus accounting for the information of the packets sent to other THEMIS stations. This information is never encrypted and can always be retrieved, even when the payload is encrypted.

from AP_i in the time $M \cdot T$ as:

$$N^i = (SN_M^i[\text{Last}] - SN_1^i[\text{First}]) \bmod 4095.$$

Let us also denote $E[L_i]$ the average bit length per packet at IP layer over *all* the packets received by station k when it is connected to AP_i . We make the reasonable hypothesis that $E[L_i]$ does not change between the connection and disconnection time from AP_i . Finally, we calculate the AP_i backhaul utilization rate as

$$\beta_i = \frac{E[L_i] \cdot N^i}{M \cdot T}. \quad (9)$$

Wireless Capacity

THEMIS measures the wireless capacity by calculating the packet dispersion of frames directed to it when the AP is transmitting in saturation. In order to detect saturation periods, station k run-time senses the *wireless channel occupancy*, that is, the percentage of time that the channel is busy, between two consecutive received packets. These statistics are collected from specific 802.11 baseband registers, exposed by the NIC card. If the occupancy is above a certain threshold (80% in our implementation), we define the AP in saturation for that pair and store the packet length of the second packet and the dispersion between the packets. Then, referring to Fig. 5, w_{ik} is derived averaging over the window of measure $M \cdot T$ as

$$w_{ik} = \frac{\sum_{j=1}^M B_j}{\sum_{j=1}^M T_{j,SAT}^i}, \quad (10)$$

where B_j is the sum of the packet length in saturation sent from AP_i to station k and $T_{j,SAT}^i$ is the sum of the dispersions when station k receives in saturation mode during the j -th connection to AP_i . Note that w_{ik} takes into account the existing interference, and depends on the current PHY rate of APs and stations, the signal quality, and the performance anomaly [9] during the measurement period.

AP Backhaul Capacity

Several Internet services can be used to estimate the AP backhaul capacity b_i ¹⁰, some of them also provided by ISPs to their clients. Usually, a file coupled to a script is downloaded from a server. The script detects when the client has completed the download and determines b_i .

The server report may be hindered by the cross-traffic rate of the packets (eventually) being sent through the same AP_i backhaul to the other stations. THEMIS connects to a capacity server, but instead of relying on the server report, it calculates the peak reached by the utilization rate β_i during the connection time to the capacity server as

$$b_i = \max_{l=1,2,\dots,L} \overline{\beta_i[l]},$$

where L represents the number of measures during the test at the $1/(M \cdot T)$ rate, and $\overline{\beta_i[l]}$ denotes the smoothed average of $\beta_i[l]$ after the l -th calculation.

3.3 Multiple APs Manager

In order to provide an *efficient* TDMA implementation in THEMIS, the wireless driver on top of the single radio interface is *virtualized*, i.e., it appears as independent Virtual

¹⁰See for example <http://www.bandwidthplace.com> or <http://www.speedtest.net>.

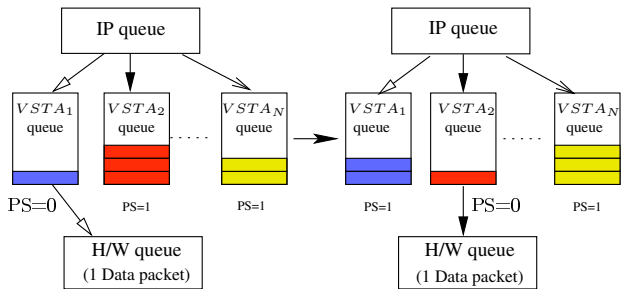


Figure 6: Queue Management.

STations ($VSTA_i$) associated to their respective APs. Each $VSTA_i$ is responsible for managing the data communication with AP_i and the related procedures such as association, authentication and scanning. To prevent losses during the TDMA operation, each THEMIS station k uses the 802.11 Power Save (PS) feature as follows (Fig. 6):

- During the active *duty cycle*, $VSTA_1$ exchanges traffic according to the 802.11 protocol, while the other $VSTAs$ are dormant in PS mode. During the PS time, both the AP_1 and the station can only buffer packets [2, 5, 15].
- When the *duty cycle* expires, $VSTA_1$ sends a frame to inform AP_1 that it is going into PS mode. Once received the MAC ACK, $VSTA_1$ and AP_1 start to buffer the packets destined to each other.
- THEMIS assigns the control of the card to $VSTA_2$ and switches to the AP_2 radio-frequency.
- $VSTA_2$ sends a frame to AP_2 to indicate that it is ready to send/receive traffic, and awaits for the MAC ACK.
- The process continues until the station has cycled through all the $VSTAs$. The spare duty cycle can be used for other operations such as scanning or sleeping (see Fig. 4). The station then restarts the TDMA cycle.

In order to minimize the switching cost c_i in (8), THEMIS achieves a fine-grained timing at MAC/PHY level, using the following techniques:

- THEMIS introduces a *MAC virtual queue* per AP. This allows to buffer packets in the MAC virtual queue, when THEMIS is selecting some other AP.
- THEMIS efficiently manages a hardware buffer (common to *all* the $VSTAs$) of one (1) data packet to quickly switch among MAC virtual queues. This is a challenging task, because short H/W queues cause inefficiencies that negatively affect throughput (as a comparison, the original MadWiFi driver sets the H/W queue size to 200)¹¹.
- In order to switch the PS state, THEMIS piggybacks the MAC PS bit on the header of the pending data on top of the MAC virtual queue. THEMIS reverts to the classical use of probes (as done in [2, 5]) in the rare event of not having data packets ready for transmission.

¹¹Packets in the hardware queue must be sent before the end of the *duty cycle* assigned to the $VSTA$. This causes a delay respect to the expected end of the duty cycle imposed by the THEMIS scheduler. The efficient management of a hardware buffer of size one minimizes any extra-delay.

With the techniques described above, THEMIS incurs in a switching cost c_i of about 1.2-1.5 ms, most of which (around 800 μ sec) is spent in hardware radio-channel commutation. This limited overhead, significantly less than [2, 5], increases the stability of the system by reducing the jitter in the switching procedure. This enables a fine-grained selection of duty cycles assigned by the scheduler even if the station transmits in saturation mode, which is of particular importance for TCP traffic.

On top of the MAC implementation, THEMIS uses a *flow mapper* to assign new TCP flows from the upper layers to a specific $VSTA$. While we could use a more sophisticated flow mapper, we employed a proportional based mapper as in [2]: the amount of traffic r_{ik} assigned to AP_i maintains the proportions of the bandwidth obtainable from each AP and equal to $r_{ik} = \frac{f_{ik} w_{ik}}{\sum_j f_{jk} w_{jk}}$.

Finally, THEMIS implements a Reverse-NAT module that i) makes sure that the packets leave the station with the correct source IP address (i.e. the one corresponding to the outgoing $VSTA$, as assigned by the AP); and ii) presents a consistent (dummy) IP address to the applications, providing IP transparency to higher layers.

4. VALIDATION

We evaluate THEMIS in an extensive set of tests. Our findings show that

- the estimators described in Section 3.2 are accurate, and stations do not need to request information from the network.
- THEMIS achieves a fair sharing of the aggregate network capacity among stations, while efficiently using the aggregated network capacity.

In our experiments, the APs are off-the-shelf Linksys, running Linux DD-WRTv24 firmware. The stations are Linux laptops, equipped with a single-radio Atheros-based wireless NIC. For every AP and station in the network, the wireless multimedia extensions (WME) and the RTS/CTS handshake are disabled. Any non-standard compliant 802.11 feature is also disabled, and H/W queues are set up with 802.11 best effort parameters.

4.1 Evaluation of THEMIS Estimators

We first verify the accuracy of the estimators used by THEMIS. We start studying the estimation of the backhaul utilization rate β_i in a test where 3 THEMIS stations download HTTP files using 3 Mbps lines. Stations are connected to the AP using a fixed connection time of 25 ms over a period of 100 ms. Stations are not synchronized, and they connect to the corresponding APs at independent times. Consequently, stations can only observe a fraction of the traffic load sent to other stations. Moreover, because of the wireless nature, they may not receive some packets sent to other stations, missing information such as the sequence number SN and packet length L_i needed by the estimator in (9).

In this configuration, we compare the estimations of the backhaul utilization rate over the time at each THEMIS station with the actual rate measured at the AP. The results in Fig. 7 show that all stations obtain a very accurate estimation.

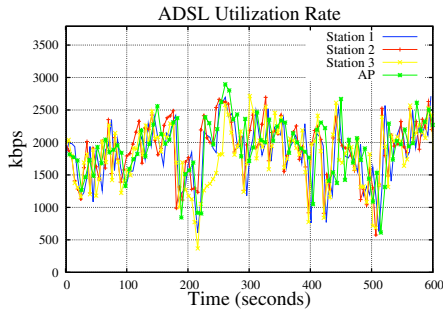


Figure 7: AP backhaul utilization estimation (β_i).

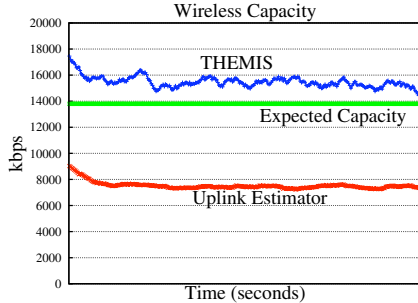


Figure 8: Wireless capacity estimation (w_{ik}).

We next evaluate the THEMIS’ wireless capacity estimator described in Section 3.2. In the test, the THEMIS station connects to an AP with a duty cycle of 25 ms over a period of 100 ms, and performs several HTTP downloads from different Internet servers. Fig. 8 shows the estimation of w_{ik} in a period of 4 minutes. THEMIS estimator gives a good approximation (around 13.7 Mbps) of the speed reported with a downlink Iperf test from a server located in the same LAN of the AP.

Estimators of w_{ik} are also proposed in [2, 16]. However, these estimators are based on the time needed to transmit a packet from the 802.11 station, and so they better represent uplink speeds rather than downlink. This can result in severe errors in the estimation of the downlink wireless capacity. As an example, Fig. 8 shows the performance of the estimator in [2] for the same scenario, and we observe that it under-estimates the wireless capacity. In fact, a high downlink speed will cause a long air-time before transmitting a packet in uplink, that translates in a low (and erroneous) downlink wireless capacity estimation.

4.2 System Evaluation

We now evaluate the system implementation of THEMIS through different tests. For every scenario, we run five tests of 1800 secs and plot the average results obtained. We choose such a configuration to verify that results are stable in time and across different tests. To achieve independent tests, stations are configured so that the THEMIS estimators are reset after each test. For the transport layer, we use Linux standard TCP Reno with SACK and delayed ACK option enabled and we emulate the AP backhaul capacities using the *tc* Linux traffic shaper. Unless otherwise stated: i) we open a TCP flow per AP using *iperf*, ii) the AP backhaul capacity is known at each station k , while the ADSLs

	station A	station B
802.11 Legacy	0.45 Mbps	6.24 Mbps
THEMIS	3.15 Mbps ($f=0.19$)	3.40 Mbps ($f=0.15$)

Table 1: Two stations connected to one AP.

utilization rates $\{\beta_{ik}\}$, and the wireless capacities $\{w_{ik}\}$ are estimated at THEMIS station k as described in Section 3.2.

THEMIS Parameters

Selecting the appropriate wireless period represents a complex trade-off. On one side, switching among APs introduces overhead, so selecting long *wireless periods* is more efficient. However, long periods affect TCP performance because they artificially increase the end-to-end delay. On the other hand, short periods reduce the disconnection time from the APs in PS mode, and prevent TCP from timing-out, but are more inefficient. As a good balance, we select a *wireless period* T of 100 ms. The scheduler and estimators are updated every $20 \cdot T = 2$ seconds. We also impose that the time of connection to each AP is at least equal to the switching cost plus 2 ms (that gives a minimum *duty cycle* $f_i \geq 0.03$).

The values of α (5), δ (6) and γ (7) have been selected based on extensive simulations, with values that provide a good trade-off between convergence and stability. Similarly, we choose the congestion thresholds for the AP backhaul and the wireless capacity to be $\lambda=0.95$ and $\mu=0.95$ respectively. A more detailed sensitivity analysis of the parameters falls outside the scope of this paper.

Two Stations Connected to One AP

We first consider the configuration where two stations are connected to the same AP (802.11 legacy operation). In the test, we consider that both stations receive traffic from the AP at a downlink wireless rate of about $w_1=20$ -22 Mbps and are connected to an AP backhaul of $b_1=7$ Mbps¹². We also consider that station A opens one TCP flow per AP while station B opens 10 TCP flows per AP.

The results are summarized in Table 1. With legacy 802.11, station B uses most of the backhaul capacity with an average received throughput of 6.24 Mbps while station A starves at 0.45 Mbps, at a throughput more than 13 times smaller than station A. On the other hand, each THEMIS station connects for a limited percentage of card time on each AP to collect the requested bandwidth T_{1k} . The result is that station B — that opens more flows — connects less time than station A, i.e. 14% versus 19% of their time, and then for just a few ms of the entire wireless period. Indeed station B needs in average less time to achieve the bandwidth from the AP, because it is less affected by the TCP’s sawtooth behavior of each flow. As a result, stations A and B obtain similar throughput (3.15 Mbps vs 3.40 Mbps), with a network utilization of 6.55 Mbps instead of 6.69 Mbps, a consequence of the THEMIS congestion control.

¹²This is the AP backhaul capacity, and hence the actual speed available for TCP traffic may be lower. In fact, because of TCP’s sawtooth behaviour, not all the *available bandwidth* at the bottleneck may be used at any time. The bandwidth utilization per path can increase establishing more than one TCP connection over each AP.




Figure 9: Duty cycles evolution with one station.

	AP_1	AP_2	
Capacity	$b_1=5$ Mbps	$b_2=5$ Mbps	$b_2=2.73$ Mbps
Duty Cycle	0.25	0.71	0.67
Throughput	4.74 Mbps	2.43 Mbps	2.21 Mbps

Table 2: Connection to two APs, one wireless bottleneck.

One Station Connected to Multiple APs

In these tests we evaluate the efficiency in terms of network utilization with one THEMIS station connected to multiple APs. Let us first consider the case where the throughput is not limited by the wireless card speed on any of the connections, i.e. the expected result is to completely utilize the available backhaul capacity of the APs.

Consider station A is associated to 3 APs, at a downlink wireless rate of about $w_1=w_2=w_3=20$ Mbps and is connected to AP backhauls of $b_1=5$ Mbps, $b_2=1$ Mbps and $b_3=10$ Mbps respectively, with a total bandwidth of 16 Mbps. As we can see in Fig. 9, the duty cycles converge to stable range of values. THEMIS spends most of the time on the best network path (via AP_3) and less time on the worst network path (via AP_2). This results in a total aggregated throughput of 15.05 Mbps, that is with an average utilization of 94% of the network aggregated capacity, as we expect from the setting of $\lambda=0.95$.

We then consider a scenario where a THEMIS station connects to two APs, and is limited by the wireless speed on one link. In the test, a THEMIS station measures a downlink wireless capacity of $w_1=20.74$ Mbps on AP_1 and $w_2=2.73$ Mbps on AP_2 and is connected to AP backhauls of 5 Mbps each, bottlenecked by the wireless on path 2.

Results are summarized in Table 2. We consider two settings: first, the ideal case where the AP backhaul capacities are correctly estimated at 5 Mbps, and second, the most realistic scenario where the estimation of the AP backhaul capacity of the path limited by wireless (path 2) is bottlenecked by the wireless capacity $b_2=w_2=2.73$ Mbps.

In the first case, THEMIS spends $f_1=0.25$ on the path with higher wireless speed, obtaining a throughput of 4.74 Mbps. The rest of the time it is spent in the path limited by the wireless link ($f_2=0.71$), where it achieves a throughput of 2.43 Mbps (for an aggregated 7.17 Mbps). Note that a small time ($f=1-0.25-0.71=0.04$) is used by THEMIS to detect card time congestions as shown in (7).

In the second case, the throughput achieved on path 2 slightly reduces to 2.21 Mbps, with a sub-utilization of the path of $2.43-2.21=0.21$ Mbps. In fact, a smaller (and wrong)

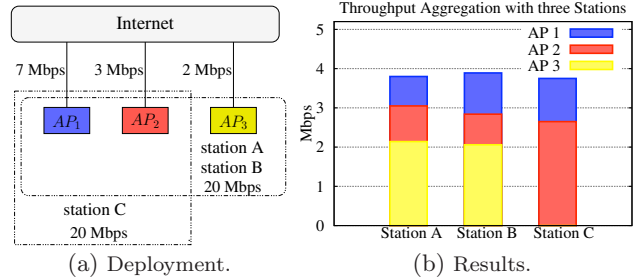


Figure 10: 3 THEMIS stations sharing 3 APs.

AP backhaul capacity estimation causes a higher AP backhaul price p_2 on the link, that in turn causes the station to request less throughput on this connection according to (5). This translates in a smaller duty cycle $f_2=0.67$ rather than 0.74, that in turns reduces the bandwidth received on this path.

In both tests, THEMIS makes an efficient usage of the network: the overall throughput is higher than the one obtained being connected 100% of the time to AP_1 (at most 5 Mbps) or to AP_2 (at most $w_2=2.73$ Mbps).

Multiple Stations Connected to Multiple APs

We evaluate the fairness and the network utilization efficiency, when different stations are connected to multiple APs. First, we analyze the case of 3 THEMIS stations, in the scenario in Fig. 10(a), with 3 APs with backhaul speeds of $b_1=7$ Mbps, $b_2=3$ Mbps and $b_3=2$ Mbps respectively, resulting in a total aggregated capacity of 12 Mbps. Given that none of the stations is limited by the wireless links, each station is expected to get an average aggregated speed close to $12/3=4$ Mbps, even if the stations share a different number of APs.

This is demonstrated in Fig. 10(b): the 3 stations obtain a fair share of the aggregate AP backhaul speed, averaging 3.80 Mbps, 3.89 Mbps and 3.75 Mbps on stations A, B and C, respectively, for a total aggregate throughput of 11.44 Mbps, again around the 95% of the overall available capacity.

Then we consider the scenario in Fig. 11(a), where station B shares two AP backhauls with station A at wired speeds of 5 and 1 Mbps, respectively. Station A can also connect to a third AP (AP_3) with a backhaul speed of 10 Mbps. Then, station B can obtain at most 6 Mbps and can never reach the 10 Mbps speed of AP_3 backhaul.

The results in Fig. 11(b) show a total aggregate TCP throughput of 9.88 Mbps on station A (with $f_1=0.08$, $f_2=0.05$ and $f_3=0.47$), and 5.09 Mbps on station B ($f_1=0.28$, $f_2=0.09$). Station A makes the fair decision, reducing the amount of time connected to the shared APs as much as possible.

Stations With an Uneven Number of TCP Flows

Let us recall the flow distribution unfairness example shown in Fig. 1(a) (Section 2) where two stations are sharing two 5 Mbps backhaul APs and use an uneven number of TCP flows. Fig. 12 shows that THEMIS is able to guarantee a fair share of the aggregated backhaul capacity to each station.

Stations With Different Priorities

Consider the same scenario as before, where now station A and station B happen to be roaming and sharing two 5

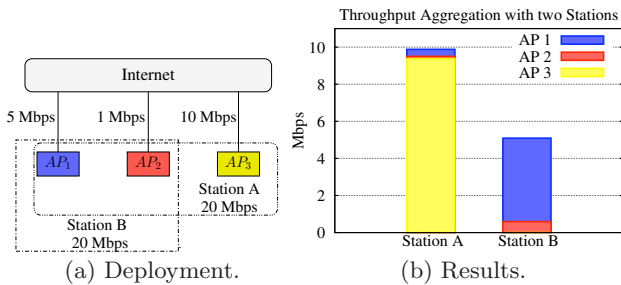


Figure 11: Two stations sharing partially overlapping sets of APs where station B cannot obtain the throughput obtained by station A.

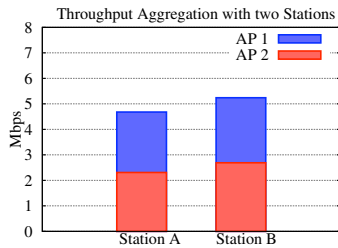


Figure 12: Fair share of the backhaul bandwidth using THEMIS.

Mbps AP backhauls. Let us consider that station A belongs to a user that has higher priority than the user of station B. For sake of illustration, we suppose that THEMIS applies weighted proportional fairness using $K_A = 4$ and $K_B = 1$. Therefore, it is expected that station A obtains $K_A/(K_A + K_B) = 0.8$ of the total bandwidth while station B obtains the remaining $K_B/(K_A + K_B) = 0.2$. The experiments show that THEMIS stations obtain a throughput of 7.64 Mbps for station A and 2.0 Mbps for station B.

5. THEMIS IN THE WILD

In order to test the scalability of THEMIS, we deploy a realistic testbed spanning three floors of a multistory building. The network consists of 10 commercial ADSLs with their corresponding WLAN APs and 10 THEMIS stations, i.e. the owners of each line. Nine of the ADSL lines have a nominal capacity of 3 Mbps and one has a nominal capacity of 1 Mbps. The APs are distributed every 80 square meters to emulate the average residential flat size (see Fig. 13) and are set to independent radio-frequencies in the 2.4 GHz ISM band¹³.

In the bootstrap phase, the APs are selected based on a passive analysis of the SNRs of the 802.11 AP beacons. Stations scan for the APs in range and start authenticating and associating to the APs, starting with the ones with highest SNR down to the ones with smaller SNR. THEMIS requires a minimum SNR of 10dB to guarantee a stable reception at 1 Mbps PHY basic rate. In each test, automatic rate selection is active in each THEMIS station, with independent instances of the Minstrel rate selection algorithm [17] over each wireless uplink.

¹³The channels optimization is out-of-the-scope of this paper.



Figure 13: Testbed deployment. APs and stations have been deployed over 3 floors, ground floor (on the left), mezzanine (in the middle), and first floor (on the right). Each circle represents an AP, while stations are placed nearby the APs, one station per AP. Only stations A, B and C, relevant for some experiments, are shown in the map. Obstacles, as walls and desks are presented between all the AP links.

5.1 Characterization

We measure the capacity of each link of the network (i.e. the ADSLs and the 10×10 wireless links). Our findings are that the 3 Mbps lines offer a constant maximum speed of 2.65 Mbps and the 1 Mbps line offers 0.89 Mbps. Regarding the wireless links, apart from the 10 “home” links where the station is located nearby the AP, the SNR measured per wireless link is consistently lower than 30 dB.

We then generate traffic from a server connected to the APs via an 802.3 LAN, activating one AP-station link at a time, with 5 minutes dedicated to each test. We calculate the average throughput and the standard deviation for each link. Then, we re-order the 10 links in descending order per-station, based on the average throughput.

Results are reported in Fig. 14. Each station can receive TCP traffic from at least 3 APs (and up to 5) at a speed higher than 10 Mbps. The results show the feasibility of aggregating the low-speed backhaul bandwidth of at least three APs.

5.2 The Effect of Location

To show the effect of location, we perform a test in which two stations (station A and station B as shown in Fig. 13) initially share the same set of APs and are located a few meters away from its “home” AP. For this test we use three APs connected to 3 Mbps lines, hence, the total backhaul capacity that station A and station B share is $2.65 \times 3 = 7.95$ Mbps. As a result, a fair share of the total bandwidth would be $7.95/2 = 3.975$ Mbps per station. Both stations perform several HTTP downloads per AP during 2400 seconds. After 1200 seconds of test, station B moves to a second location from which it can only be connected to two of the former APs. As we do not implement IP mobility in our testbed, all the connections of station B are dropped and started again in the new location. As a consequence of the movement of station B, the topology of the network changes and stations observe an uneven AP backhaul capacity.

We run the test using a throughput maximization algorithm as in FatVAP [2]¹⁴ (Fig. 15(a)), and using THEMIS

¹⁴We implemented the throughput maximization algorithm according to the description in [2]. To provide a fair compar-

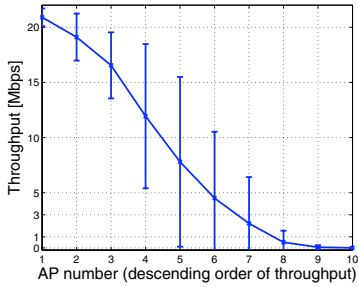


Figure 14: Wireless link quality assessment.

(Fig. 15(b)). Results show that, when the network topology is similar for both stations (they are both connected to 3 APs at similar speed), using throughput maximization results in a similar long-term performance for both stations, but with no guarantee of short-term fairness. Moreover, when the topology changes, station B is clearly penalized by its new unlucky location obtaining 2.8 Mbps while station A obtains 4.8 Mbps.

On the other hand, THEMIS guarantees a fair share of the backhaul capacity in both topologies, offering 3.5 Mbps to each station. Note that when station B moves to the new position, the PHY rate is quickly reduced because of the lower signal strength, with THEMIS quickly converging to a fair assignment of the backhaul capacity. Also note that because the fairness mechanism relies on the congestion thresholds λ and μ (Section 2.2), the network utilization is slightly lower than the optimal.


5.3 Integrated Operations

We have shown via different deployments that THEMIS is able to deal with the three types of unfairness that arise when aggregating AP backhaul bandwidth. However, in a real life scenario, these unfairness can take place at the same time. Thus, we perform a test that evaluates THEMIS in presence of a P2P station (station A), an unluckily located station (station B) and a low priority station (station C). The location of the stations is shown in Fig. 13. For this test we use three APs, each with a 3 Mbps backhaul. The P2P and the low priority stations are connected to 3 APs while the unluckily located station is connected to 2 APs. Given that the low priority station owns a 1 Mbps ADSL while the others own a 3 Mbps ADSL line, the weights have been set to $K_A = K_B = 3$ and $K_C = 1$. In such experiment, a fair system should be able to allocate the bandwidth proportionally to the priority of the users.


At the beginning of the test, station A starts downloading P2P traffic from the three APs. After 1200 seconds, station B starts several HTTP downloads from the two APs it is connected to. Finally after 1200 seconds more, station C also starts HTTP traffic from the APs.

The result of using a throughput maximization algorithm is shown in Fig. 16(a). It is noticeable that station A, due to the high number of TCP flows opened by P2P applications, obtains most of the backhaul capacity preventing station B from obtaining its fair share of the bandwidth. Furthermore, when station C starts its downloads, the absence of prior-

ity among users further reduces the throughput obtained by station B, introducing billing unfairness. Finally, since station B and C achieve a similar throughput despite that station B is unluckily located, the flow distribution unfairness dominates over the topology unfairness.



(a) FatVAP



(b) THEMIS

Figure 15: Assessment of the topology unfairness in the residential-like deployment.


ity among users further reduces the throughput obtained by station B, introducing billing unfairness. Finally, since station B and C achieve a similar throughput despite that station B is unluckily located, the flow distribution unfairness dominates over the topology unfairness.

The result of using THEMIS is shown in Fig. 16(b). When the unluckily located station B starts its downloads after 1200 seconds, the wireless capacity measured at station A over the shared APs is reduced because of the performance anomaly [9]. However, the system quickly adapts: the wireless links with lower wireless capacity receive a higher wireless price q_{ik} and hence smaller throughput demand T_{ik} and dedicated duty cycle f_{ik} . A smaller duty cycle for both stations A and B means that the probability of being connected to the same AP at the same time, and consequently the occurrence of performance anomaly, is reduced. Concluding, THEMIS offers a fair share of the aggregated bandwidth to both stations, while providing a high usage of the backhaul bandwidth. Finally, when station C starts its downloads, the priorities are preserved and stations A and B obtain a greater share of the backhaul capacity.


6. RELATED WORK

In recent years Wi-Fi communities have attracted the attention of both the research community and the wireless industry because of the uptake of WLAN in residential areas. In this direction [3, 4, 18] propose to allow members of the communities to share the backhaul bandwidth of their WLAN APs. Among those, Wi-Sh [4] discusses the fairness problems that can arise from sharing resources. However, it does not consider the use of multiple APs to aggregate their backhaul bandwidth.

Backhaul bandwidth aggregation has been explored in



(a) FatVAP



(b) THEMIS

Figure 16: Test with the effect of the three types of unfairness: station A uses P2P traffic, station B is unluckily located (starts after 1200s) and station C is a low priority station (starts after 2400s).

[19, 20], where stations connect to their home APs via ethernet and to the remote APs using WLAN. However, they do not connect to multiple APs via WiFi, so the number of APs they can aggregate is limited by the number of physical interfaces (ethernet and WiFi) available at the stations.

The idea of connecting to multiple APs through a single radio was first shown in VirtualWiFi [5]. The authors rely on the WLAN standard power saving (PS) mode to switch among different Wi-Fi nodes in time division. Switching among Wi-Fi nodes is transparent to the applications, but at a high cost in time (30-600 ms). In fact, VirtualWiFi implements the code on top of the driver card with a MAC instance for connection.

Within the problem of single radio AP backhaul aggregation, the closest work is FatVAP [2]. The authors introduce a scheduler to select the percentage of time to spend on each AP to maximize the aggregate throughput at each station. However [2] has a limited focus because it does not resolve the unfairness across stations, and it only considers stations connected to (strictly) more than one AP. Furthermore, the local throughput maximization approach in [2] can not be extended in order to take into account priority-based per-station fairness. Compared to [2], THEMIS fairly aggregates the AP backhaul bandwidth among the different THEMIS stations, irrespectively of their location, link quality and number of APs they have in range. Moreover, THEMIS is able to adapt to different fairness objectives in order to accommodate the different scenarios discussed in Section 1, and it achieves this in a completely distributed manner. Finally, THEMIS implementation of the single-radio multi-AP TDMA access is improved compared to [2, 5], reducing the

frequency switching overhead and increasing the accuracy when selecting the amount of time that the station connects to the different APs. This results in a more efficient operation and increased throughput.

Among other work, [21] introduces a support for a seamless hand-off between WLAN APs. In [15], standard solutions have been exploited to increase the aggregate throughput observed by a single station with respect to the design in [2, 5, 21]. However, these works do not consider the problem of fairness.

Link-alike [22] tackles the problem of minimizing the uplink total transfer time via multiple wireless links. However, the solution requires cooperation among the APs, with 802.11 APs transmitting and receiving at the same radio-frequency, and a custom TCP protocol over the wireless link.

Several tools have been designed to estimate the available bandwidth along a network path. However, these tools typically send active probes along a path and/or require a cooperative implementation at both the sender and receiver [23, 24].

There is little work on non-cooperative estimation of the ADSL available bandwidth. Most notably ABwProbe [25] and FAB-Probe [26] rely on the asymmetry of the ADSL downlink capacity to send TCP ACK packets of different sizes and receive small TCP RST packets from the TCP client. Since the TCP RST is at fixed length, they cannot estimate the available bandwidth from the client-side, as done by THEMIS.

The estimation of the wireless capacity has been studied with different levels of accuracy (see e.g. [2, 27]). A comparison with the implementation of THEMIS has been provided in Section 4.1. Our experimental evaluation has demonstrated the robustness of THEMIS in realistic scenarios, under MAC contention, adaptive PHY rates and performance anomaly.

7. CONCLUSION AND FUTURE WORK

We have shown that fairness is a crucial factor for the success of multi-AP aggregation schemes. Without fairness, the perceived value of the system is severely reduced, eliminating the incentives of users to participate, and of providers to deploy it. This effectively renders the scheme unfeasible. In order to achieve fairness, existing multi-AP aggregation systems that maximize the throughput of single users cannot be extended. As a consequence a complete re-design of the system is required.

To address this problem we introduced THEMIS, a single-radio station implemented in commodity hardware that is fair in a multi-AP aggregation scenario. THEMIS operates locally at the station, using standard 802.11, without requiring any change in the network. This makes THEMIS ready to be deployed. In fact, THEMIS is being used by a major broadband provider in a commercial trial.

There are several interesting lines to expand this work. On the technical side, we plan to extend THEMIS to include uplink traffic in the formulation, and investigate the impact and trade-offs that TDMA may have over the TCP performance. From an architectural point of view, we are currently exploring the use of THEMIS to design more power efficient access networks. Finally, it would be interesting to understand how THEMIS can be leveraged to perform efficient large-scale cellular data offloading, which appears to be a difficult challenge for the years to come.

8. REFERENCES

- [1] D. Han, A. Agarwala, D. Andersen, M. Kaminsky, D. Papagiannaki, and S. Seshan, "Mark-and-Sweep: Getting the 'Inside' Scoop on Neighborhood Networks," in *Proc. of the ACM IMC Conf.*, (Vouliagmeni, Greece), pp. 99–104, October 2008.
- [2] S. Kandula, K. Lin, T. Badir Khanli, and D. Katabi, "FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput," in *Proc. of the USENIX NSDI Conf.*, (San Francisco, CA), pp. 89–04, April 2008.
- [3] "FON." <http://www.fon.com/>.
- [4] 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.*, (Rio de Janeiro, Brazil), pp. 1638–1646, April 2009.
- [5] R. Chandra and P. Bahl, "MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card," in *Proc. of the IEEE INFOCOM Conf.*, (Hong Kong, China), pp. 882–893, March 2004.
- [6] E. Park, D. Kim, H. Kim, and C. Choi, "A Cross-Layer Approach for Per-Station Fairness in TCP over WLANs," *IEEE Trans. Mobile Comput.*, vol. 7, no. 7, pp. 898–911, 2008.
- [7] B. Briscoe, "Flow Rate Fairness: Dismantling a Religion," *ACM SIGCOMM Comp. Commun. Review*, vol. 37, no. 2, pp. 63–74, 2007.
- [8] G. Tan and J. Guttag, "Time-based Fairness Improves Performance in Multi-rate WLANs," in *Proc. of the USENIX Annual Tech. Conf.*, (Boston, MA), pp. 23–24, June 2004.
- [9] F. R. G. B.-S. M. Heusse and A. Duda, "Performance Anomaly of 802.11b," in *Proc. of the IEEE INFOCOM Conf.*, vol. 2, (San Francisco, CA), pp. 836–843, April 2003.
- [10] H. J. Kushner and P. A. Whiting, "Convergence of Proportional-Fair Sharing Algorithms under General Conditions," *IEEE Trans. Wireless Commun.*, vol. 3, pp. 1250–1259, 2003.
- [11] R. Srikant, *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. Springer Verlag, 2004.
- [12] M. P. W. Wang and S. H. Low, "Optimal Flow Control and Routing in Multi-path Networks," *Elsevier Perform. Eval.*, vol. 52, no. 2-3, pp. 119–132, 2003.
- [13] "Madwifi project." <http://madwifi-project.org>.
- [14] E. Kohler, R. Morris, B. Chen, J. Jannotti, and F. M. Kaashoek, "The Click Modular Router," *ACM Trans. Comput. Syst.*, vol. 18, pp. 263–297, August 2000.
- [15] D. Giustiniano, E. Goma, A. Lopez Toledo, and P. Rodriguez, "WiSwitcher: An Efficient Client for Managing Multiple APs," in *Proc. of ACM PRESTO Wrkshp.*, (Barcelona, Spain), pp. 43–48, August 2009.
- [16] D. Malone, I. Dangerfield, and D. J. Leith, "Verification of Common 802.11 MAC Model Assumptions," in *Proc. of the ACM PAM Conf.*, (Louvain-la-Neuve, Belgium), pp. 63–72, April 2007.
- [17] "Minstrel rate control algorithm." <http://linuxwireless.org/en/developers/Documentation/mac80211/RateControl/minstrel>.
- [18] "Whisher Wifi Sharing Community." <http://www.whisher.com>.
- [19] E. Tan, L. Guo, S. Chen, and X. Zhang, "CUBS: Coordinated Upload Bandwidth Sharing in Residential Networks," in *Proc. of the IEEE ICNP Conf.*, (Plainsboro, NJ), pp. 193–202, October 2009.
- [20] N. Thompson, G. He, and H. Luo, "Flow Scheduling for End-host Multihoming," in *Proc. of the IEEE INFOCOM Conf.*, (Barcelona, Spain), pp. 1–12, April 2006.
- [21] A. Nicholson, S. Wolchok, and B. Noble, "Juggler: Virtual Networks for Fun and Profit," *IEEE Trans. Mobile Comput.*, vol. 9, no. 1, pp. 31–43, 2010.
- [22] S. Jakubczak, D. G. Andersen, M. Kaminsky, D. Papagiannaki, and S. Seshan, "Link-alike: Using Wireless to Share Network Resources in a Neighborhood," *IEEE Trans. Mobile Comput.*, vol. 12, no. 4, pp. 1–14, 2008.
- [23] D. K. J. Strauss and F. Kaashoek, "A Measurement Study of Available Bandwidth Estimation Tools," in *Proc. of the ACM IMC Conf.*, (Miami Beach, FL), pp. 39–44, October 2003.
- [24] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cot, "pathChirp: Efficient Available Bandwidth Estimation for Network Paths," in *Proc. of the ACM PAM Wrkshp.*, (San Diego, CA), April 2003.
- [25] D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. W. Biersack, "Non-cooperative Available Bandwidth Estimation Towards ADSL Links," in *Proc. of the IEEE Global Internet Symp.*, (Phoenix, AZ), April 2008.
- [26] D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. W. Biersack, "Fast Available Bandwidth Sampling for ADSL Links: Rethinking the Estimation for Larger-Scale Measurements," in *Proc. of the ACM PAM Conf.*, (Seoul, South Korea), pp. 67–76, April 2009.
- [27] J. P. R. Draves and B. Zill, "Routing in Multi-radio, Multi-hop Wireless Mesh Networks," in *Proc. of the ACM MobiCom Conf.*, (New York, NY), pp. 114–128, October 2004.

APPENDIX

If after applying the correction factors in eq. (8), the resulting corrected duty cycles are such that $\sum_i f_{ik} > 1$, we apply the following algorithm to distribute the spare duty cycle:

1. we first reduce the duty cycles for those stations that overestimated it, i.e., we recalculate the adjusted duty cycles f'_{ik} as follows

$$f'_{ik} = \begin{cases} \sigma_{ik} f_{ik} & \text{if } \sigma_{ik} \leq 1 \\ f_{ik} & \text{otherwise} \end{cases}$$

2. Once adjusted, if the demanded duty cycles exceeds the capacity of the card, i.e., $\sum_i f'_{ik} > 1$, then we normalize them $f''_{ik} = f'_{ik} / \sum_i f'_{ik}$. If, on the other hand, there is spare time $f_{sp} = 1 - \sum_i f'_{ik}$, we distribute it among the links that need to increase their duty cycles ($\sigma_{ik} > 1$) proportionally to their needs as follows

$$f''_{ik} = \begin{cases} f'_{ik} + f_{sp} \frac{\sigma_{ik}}{\sum_i \sigma_{ik}} & \text{if } \sigma_{ik} > 1 \\ f'_{ik} & \text{otherwise} \end{cases}$$

3. Each station uses the resulting values f''_{ik} .