

Improving Flow level Fairness and Interactivity in WLANs using Size-based Scheduling Policies

Guillaume Urvoy-Keller
Eurecom
2229, route des crêtes
06904 Sophia-antipolis, France
guillaume.urvoy@eurecom.fr

Andre-Luc Beylot
ENSEEIH
Laboratoire IRIT
31071 Toulouse Cedex 7, France
andre-luc.beylot@enseeiht.fr

ABSTRACT

In this paper, we investigate the use of a size-based scheduling policy, LASTOTAL, in WLANs. A size-based scheduling policy is a priority policy where the priority of a flow is based on its size. LASTOTAL replaces the legacy IP level FIFO scheduler at the access point. The lower protocol layers, and especially the MAC 802.11 layer are left unchanged. We demonstrate using realistic synthetic workloads, that LASTOTAL solves the unfairness issue due to DCF in 802.11 WLANs and ensures small response times to the majority of the flows under any load conditions. The latter property is desirable as short flows correspond to interactive applications and maintaining low response times for those flows despite load variations, significantly improves user experience. We also introduce and validate Markovian queuing models to assess the response time of the access point for both FIFO and LASTOTAL.

Categories and Subject Descriptors

C.2.5 [Local and Wide-Area Networks]: Internet (e.g., TCP/IP); C.4 [Performance of Systems]: Performance attributes

General Terms

Performance

Keywords

Size-based scheduling, LAS, TCP, 802.11

1. INTRODUCTION

Size-based scheduling has proved to be very effective in increasing performance in a lot of scenarios: Web servers [6], Internet traffic [15] or 3G networks [9]. The key idea behind size-based scheduling is to favor short jobs while ensuring that large jobs do not starve. The net result is better interactivity from the user point of view and better performance

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.

MSWiM'08, October 27–31, 2008, Vancouver, BC, Canada.
Copyright 2008 ACM 978-1-60558-235-1/08/10 ...\$5.00.

as the number of active jobs at any time instant tends to be smaller, which prevents buffer overflows. The extent to which large jobs suffer depends on the statistical characteristics of the job size distribution and especially how the mass is distributed among short and large jobs. Broadly speaking, the larger the mass carried by the large flows, the smaller the penalty since short flows, that have the highest priority, can not monopolize the server. Heavy-tailed distributions, which have often been observed in the Internet [5], feature such a property.

In this paper, we investigate the use of a size-based scheduling policy, the Least Attained Service (LAS) policy [13], in a WLAN setting. We propose an approach whereby only the IP level scheduler at the access point needs to be modified, while leaving the lower protocol layers, and especially the MAC 802.11 layer unchanged. Our first contribution is to pinpoint that the shared nature of wireless medium in 802.11 networks introduces a coupling between the going and reverse path of bidirectional connections (any TCP and a lot of UDP transfers, e.g., VoIP), which requires any size-based scheduling approach to simultaneously account for both directions while scheduling packets. As a result, we propose a specific implementation of LAS, that we call LASTOTAL, where the priority of a packet is set equal to the sum of bytes carried by the corresponding connection from or to a wireless station so far.

We show that LASTOTAL has two key benefits: (i) It solves the known unfairness problem of the DCF coordination function of 802.11 when used in a infrastructure mode like a WLAN setting; (ii) From a user perspective, LASTOTAL improves interactivity, as short flows typically correspond to interactive applications. Most of the short flows have a reduced response time, even at low load values as compared to when the legacy FIFO policy is used on top of DCF.

Our last contribution is a Markovian model that enables to assess the total amount of time spent by the packets of a flow at the queue of the access point. A key problem here is the choice of the capacity offered by the access point. To tackle this issue, we rely on a statistical approach whereby the IP level capacity of the access point is derived from simulations results.

2. RELATED WORK

A host of works have tackled the issue of unfairness between uploads and downloads observed in 802.11 networks. In [17], the authors identify the roots of the problem, which is the competition between TCP ACKs of the uploads and

TCP data packets at the buffer of the access point, which is in general small, typically 30 to 50 packets. While both TCP data packets and ACKs are lost, the impact on downloads is not the same as on uploads since TCP reacts primarily to data segments losses and is quite resilient to acknowledgment losses.

The authors in [17] propose to enforce fairness by adjusting the advertised window of TCP connections. Their solution requires to passively estimate the RTT of each connection to adjust the advertised window.

In [2], the authors propose a new scheduling policy that allows transmission of bursts of frames where the burst size is a function of the channel failures experienced by a client. Their objective is to enforce short term fairness among competing TCP flows. Their technique requires to modify the 802.11 MAC protocol.

In [11], the authors propose AAP, a modified MAC layer where the access point is granted a higher priority to access the medium than the wireless stations. It is based on the observation that while the access point has in general more traffic to convey than the wireless stations due to the asymmetry of the traffic (more downloads than uploads), its probability to access the medium is the same as any other station. Increasing the AP priority allows to drain the buffer at the access point and thus to alleviate the unfairness issue between uploads and downloads.

In [8], various combinations of layer 2 and layer 3 policies are evaluated using measurements on a WLAN testbed close to saturation. At layer 2, AAP and DCF are considered while at layer 3, the authors investigate the benefits of LASACK as compared to FIFO. LASACK is based on LAS and aims at solving the unfairness issue of DCF by assigning a priority to the TCP acknowledgments from the uploads that accounts for the amount of data traffic that has been uploaded by the flow so far. LASTOTAL is an extension of LASACK in the sense that LASTOTAL does not make any assumption on the transport layer of a connection. Our work also differs from the one in [8] for the following aspects: (i) We demonstrate the necessity to account for the half duplex nature of 802.11 wireless links by showing how LASTOTAL outperforms LAS (Section 4); (ii) We highlight the benefits of LASTOTAL not only when the network is close to saturation but also at low loads (Section 6); (iii) We provide an analytical model of the response time of the access point (Section 7).

3. THE LEAST ATTAINED SERVICE POLICY

LAS is a size-based scheduling policy. It has been initially proposed and studied in the context of time-sharing computers in the late 60s [16]. Under LAS, priority is given to the job that has received the least amount of service. In case of ties, jobs share the server in a round-robin manner. A salient feature of LAS is that it has no internal parameter to tune.

In packet networks, we rely on the definition of LAS proposed in [15], where it becomes both a scheduling and a buffer management policy. The resulting policy works as follows. Upon reception, a packet is assigned a priority which is inversely proportional to the number of bytes sent so far by the corresponding connection (the first packet of a new connection thus has maximum priority). If ever the queue is

full upon the arrival of a new packet, this packet is assigned its priority, inserted in the queue and the packet with the lowest priority is discarded.

In packet networks where links are in general full duplex, LAS is applied to each direction of the path independently from each other. For the case of WLANs, where 802.11 makes the wireless channel half duplex, we propose that LAS be applied on a connection basis. Practically, this means that the priority of a packet at the access point is based on the total amount of traffic sent by the corresponding connection. For the specific case of a pure TCP acknowledgment of an upload connection, its priority will be set equal to the amount of bytes carried by the data packets on the other direction of the connection. We call LASTOTAL this scheduling policy. In Section 4, we demonstrate that LASTOTAL outperforms LAS.

The overhead of LAS and LASTOTAL as compared to FIFO is that per connection statistics must be kept with a size-based approach. We believe that it is an affordable task for an access point that should not service a large number of simultaneous connections.

4. ACCOUNTING FOR THE HALF-DUPLEX NATURE OF WLANS

In this section, we highlight the impact of accounting for the half-duplex nature of 802.11 wireless links. To do so, we evaluate FIFO, LAS and LASTOTAL for the case of long lived TCP transfers. Though less realistic than a workload consisting of a mix of short and long transfers, such a scenario eases the illustration of the fundamental difference between LASTOTAL and the original LAS policy.

The network configuration that we consider in this paper consists of 10 wired stations and 10 wireless hosts serviced by an access point. The protocol used in the wireless part is 802.11b, with RTS/CTS disabled and a MAC level bandwidth of 11 Mbits/s. The DCF access method is used. The 10 wireless stations are at the same physical distance of the access point and in line of sight of each other. The 10 wired stations are connected to a router with an output rate 10 times larger than its input rate, which means that its output queue never builds up. The bottleneck, if any, is thus the access point. The buffer size of the access point ranges between 20 and 60 kbytes, similarly to what can be found in most commercial access points. We use Qualnet 3.9.5 to obtain all simulation results. Each simulation lasts 500 seconds.

In [17], the authors underscore a fundamentally unfair behavior of FIFO ¹ in WLANs that results in the uploading flows obtaining a larger share of the network capacity than the downloading flows. Figure 1 depicts the ratio of the long term throughputs of the downloading and uploading flows for a scenario with 1 upload and 9 downloads.

We observe from Figure 1 that the original LAS policy is less fair than FIFO, as the upload achieves throughputs that are consistently two orders of magnitude larger than the ones of downloads. The explanation behind this observation is simple: the TCP acks of the uploads consistently have the highest priority at the access point and lock out the data packets from the downloads. In contrast, LASTOTAL

¹More precisely the FIFO scheduling policy combined with the droptail queue management policy. In this paper, we use the term FIFO to denote FIFO/droptail.

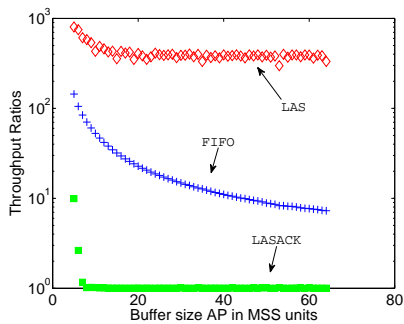


Figure 1: Ratio of upload to download throughputs

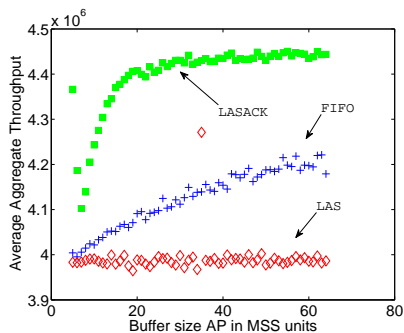


Figure 2: Aggregate throughputs

enforces a perfect fairness (ratio of 1) for all buffer sizes, as it assigns a priority to the TCP ACKs of the upload that is a function of the amount of data uploaded so far. Further note that fairness is not obtained at the expense of throughput as the aggregate throughput under LASTOTAL is larger than the ones of FIFO and LAS - see Figure 2. Given the bad performance of LAS, we consider only LASTOTAL in the remaining of this paper.

5. WORKLOAD

In a WLAN scenario with static wireless hosts, the key characteristics of the traffic are the flow arrival process and the flow duration distribution. References [7, 12] present seminal contributions in the field of WLAN traffic modeling. In those papers, several datasets corresponding to large scale WLANs comprising several hundreds APs and several thousands users are considered. In both studies, the authors observe that the distribution of flow durations and flow sizes are heavy tailed. In [7], a biPareto distribution is proposed to model the flow sizes while in [12], several candidate distributions are considered: Weibull, Lognormal, Extreme-value, Pareto.

In our simulations, we consider a Pareto distribution for the flow size. Let k be the minimum connection size and α the exponent of the power law. The density of the Pareto distribution is given by:

$$f(x) = \alpha k^\alpha x^{-\alpha-1}, \quad k \leq x, 0 \leq \alpha \leq 2. \quad (1)$$

We impose a minimum and a maximum flow size for all flows. The minimum file size is 4.5 Kbytes, i.e., 3 MSS with

an MSS value of 1.5 Kbytes. The maximum transfer size is 20 Mbytes, which constitutes a reasonable upper bound of what a user might download or upload using a WLAN access.

A key metric in the study of size-based scheduling policies is the variability of the flow size distribution. A simple metric to characterize the variability of a distribution is the coefficient of variation (CoV), which is the ratio of the standard deviation to the mean of the distribution. We can tune the coefficient of variation of a Pareto distribution through its parameter α . We consider two different values of 3 and 5.5 as CoV values, which are in the range of values commonly observed for WLAN traffic (e.g., observed values of CoV are between 2 and 6 [12]).

As for the flow inter-arrival times, the authors in [7, 12] observe that this process is stationary at small time scales (say half an hour to one hour periods). As we study the system also at small time scale (say 5 to 15 min), we freeze the parameters of our workload model for the duration of each simulation. The authors in [7, 12] observed that the flow inter-arrival process is in general more bursty than the often used Poisson process. Specifically, they propose to use a Weibull [12] or a biPareto [7] distribution to model the flow inter-arrival time distribution at small time scales, while they modulate the parameters at a larger time scales. In contrast, we rely on a Poisson process to model the flow level arrival process. This allows us to compare the simulation results with the theoretical results provided by the Markovian analysis in Section 7. In addition, while the Poisson process is less bursty than a renewal process with Weibull inter-arrival times, we can still rely on increasing the load to stress the system and observe the performance of the different scheduling strategies.

6. RESULTS

In this section, we investigate the performance of FIFO and LASTOTAL for a traffic workload consisting of a mix of flows of different sizes. We denote by λ_u and λ_d the arrival rates of TCP downloads and uploads respectively. We consider two scenarios corresponding to two values of $\frac{\lambda_u}{\lambda_d}$. For the first scenario, $\frac{\lambda_u}{\lambda_d} = 0.5$. Since we assume that the distributions of file size are similar for the two traffic directions, the upload traffic represents 33.3% of the whole traffic while the download traffic amounts for 66.6% of the total traffic. In the second scenario, $\frac{\lambda_u}{\lambda_d} = 1$, which means that the traffic workload is evenly distributed between the two directions. The latter scenario is less realistic but enables to illustrate the performance of LASTOTAL when the unfairness problem of DCF is magnified.

6.1 Aggregate Rates

We first focus on the aggregate rates for uploads and downloads under the two scheduling policies. Results are qualitatively the same for a CoV of 3 and a CoV of 5.5. We present results for the case of a CoV of 5.5. In Figures 3 and 4, we present bar plots of measured load for $\frac{\lambda_u}{\lambda_d} = 0.5$ and $\frac{\lambda_u}{\lambda_d} = 1$ where FIFO experiments are on the left and LASTOTAL experiments on the right. Results are function of the input (or offered) load, which is the product of the arrival rate to the average size of the data transfers that we use as inputs for the simulations. We vary the input load from 1 to 7 Mbits/s. As transfers are controlled by TCP

and losses can occur at the access point, the measured load can be smaller than the offered one. We distinguish for each experiment between unique (first transmission) and retransmitted bytes in both directions. We do observe that while FIFO achieves higher aggregate rates, the rates of unique data is almost equivalent under FIFO and LASTOTAL for each offered load. Under FIFO, retransmissions exclusively affect downloads, in line with the unfairness behavior of DCF/FIFO. Under LASTOTAL, retransmissions can be incurred by both uploads and downloads, proportionally to $\frac{\lambda_u}{\lambda_d}$.

The smaller retransmission ratios observed with LASTOTAL underscore the ability of LASTOTAL to enable TCP to smoothly adapt to the network resources (here the low buffer capacity at the access point). In contrast, FIFO lets TCP exhaust the network resources and do a lot of retransmissions.




Figure 3: Average load during simulation




Figure 4: Average load during simulation

6.2 Conditional Response Time

We now focus on the performance of FIFO and LASTOTAL in terms of the level of interactivity they provide. The metric we consider is the conditional response time, i.e., the average response time per transfer size (from connection set up to tear down). When presenting results, we account for the predominance of short flows over large flows. We thus present the conditional response times with respect to the percentiles of the flow size distribution. The x -th percentile of a flow size distribution is the flow size q_x such that $x\%$ of the flows have a size smaller than q_x .

The comparison between FIFO and LASTOTAL spans over several dimensions: the ratio $\frac{\lambda_u}{\lambda_d}$ of upload to download

traffic, the offered load and the variability (measured by the CoV) of the flow size distribution.

6.2.1 Offered load and traffic ratio.

Let us first consider the first two dimensions, namely the value of $\frac{\lambda_u}{\lambda_d}$ and the offered load. We consider here a Pareto distribution with a CoV of 3. Figures 5 and 7 (resp. Figures 6 and 8) present the conditional average response time for input loads of 3 and 5 Mbits/s and $\frac{\lambda_u}{\lambda_d} = 0.5$ (resp. $\frac{\lambda_u}{\lambda_d} = 1$) respectively. We observe that the offered load has more impact than the asymmetry of the traffic. Concerning the relative performance of downloads and uploads under FIFO and LASTOTAL, FIFO favors uploads at the expense of downloads. This is a new illustration of the unfairness problem of FIFO. In contrast, LASTOTAL offers almost the same service to uploads and downloads of the same size, irrespectively of the load or level of asymmetry.

We further observe that LASTOTAL does a good job at keeping the response time low for most of the uploading and downloading flows as compared to FIFO. To tune response times, LASTOTAL controls the round trip times of the connections, as can be observed in the right graph of Figures 5 to 8.

Eventually, we observe that LASTOTAL is able to maintain low response times for the short flows (i.e., a good interactivity) even when the offered load increases: the short download flows under FIFO experience a 100 fold increase in their response time when the offered load increases from 3 to 5 Mbits/s, as compared to a factor of 2 under LASTOTAL.

6.2.2 Variability of the flow size distribution.

Figures 9 and 10 present the conditional response time obtained for an input load of 5 Mbits/s for a Pareto distribution with a CoV of 5.5. When comparing those figures with the equivalent ones for a Pareto distribution with a CoV of 3 (Figures 5 and 6), we make two key observations:

- As the CoV increases, LASTOTAL offers a low response time to a larger fraction (i.e., higher percentile) of flows, in line with theory [13].
- The difference between LASTOTAL and FIFO is more pronounced for a smaller CoV, which seems counter-intuitive. Note that this phenomenon persists at offered load and $\frac{\lambda_u}{\lambda_d}$ values. We suspect that it is related to the arrival rate of fresh flows, which is higher for a CoV of 3 than for a CoV of 5.5 for the same offered load (because the average flow size is higher for a CoV of 5.5 than for a CoV of 3: 9.7 MSS against 6.7 MSS). This higher arrival rate might generate a larger level of occupancy for the buffer of the access point. One might argue that with a CoV of 5.5, even if the flow arrival rate is smaller, those flows tend to be larger. However, as they are controlled by TCP, this might explain why the overall burstiness is larger for the case CoV=3 than CoV=5.5, as large TCP transfers are controlled by the congestion avoidance algorithm and are less aggressive than short TCP transfers.

7. MARKOVIAN MODEL OF THE QUEUING TIME AT THE ACCESS POINT

In this section, we present a Markovian model to predict the response time at the access point for TCP flows in a




Figure 5: CoV=3, Input load=3 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 0.5$




Figure 6: CoV=3, Input load=3 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 1$

WLAN, i.e., the total time spent by the packets of a connection at the access point. Note that this is not the same as the response time of the flow which is the time elapsed between the sending of the first packet and of the last packet of a connection.

To the best of our knowledge, no work has tackled the issue of testing the accuracy of queuing models to estimate the response time at the queue of the access point for a workload consisting of short and long TCP transfers. The seminal work by Bianchi [1] focuses on the maximum throughput achievable under saturation, i.e., when wireless stations always have a packet to send. Bruno et al. have focused on the case of persistent TCP traffic [3, 4].

7.1 Queuing models

For wired networks, the M/G/1/PS queuing model has been shown to capture the dynamic of TCP flows with homogeneous RTTs sharing a FIFO router, while the M/G/1/LAS queue turns out to accurately model the behavior of a LAS router [14]. Application of those models for WLANs bears a number of assumptions/challenges. First, we consider single server queuing models while multiple queues can build up, either at the access point or the wireless stations. Second, we consider infinite queues, while the buffer at the access point is small.

Let the average flow arrival rate be λ . Let $f(x)$ be the probability density function of the service requirement G , $F(x) = \int_0^x f(t)dt$, its cumulative distribution function, and $F^c(x) \triangleq 1 - F(x)$. We define $m_n(x)$ as $m_n(x) \triangleq \int_0^x t^n f(t)dt$. Thus $m_1 \triangleq m_1(\infty)$ is the mean and $m_2 \triangleq m_2(\infty)$ is the




Figure 7: CoV=3, Input load=5 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 0.5$




Figure 8: CoV=3, Input load=5 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 1$

second moment of G . The load of flows with size less than or equal to x is given as $\rho(x) \triangleq \lambda \int_0^x t f(t)dt$, and $\rho \triangleq \rho(\infty)$ is the total load in the system.

We are interested in the conditional response time $E[T|G = x]$ ($T(x)$ in short), i.e., the individual response time for each flow size x . The expression for the conditional response time for the M/G/1/PS queue is:

$$T(x)_{PS} = \frac{x}{1 - \rho} \quad (2)$$

Conditional response time for an M/G/1/LAS queue is given by the following formula ([10] p. 172):

$$T(x)_{LAS} = \frac{\lambda(m_2(x) + x^2 F^c(x))}{2(1 - \rho(x) - \lambda x F^c(x))^2} + \frac{x}{1 - \rho(x) - \lambda x F^c(x)} \quad (3)$$

The LASTOTAL policy differs from the legacy LAS policy whenever upload and download data streams compete with each other to access the wireless medium. We assume that the distributions of both uploads and downloads service requirements are similar.

Under LASTOTAL, priority of a TCP ack of an upload data stream is set equal to the number of bytes sent by the corresponding data stream. Thus, in the LASTOTAL

²The distribution G accounts for the service requirement (in second) of the client TCP connections which is related to the flow size (in bytes) X by $G = X/C$, where C is the capacity of the server. Note that we assign a specific distribution - Pareto - to X and not to G .

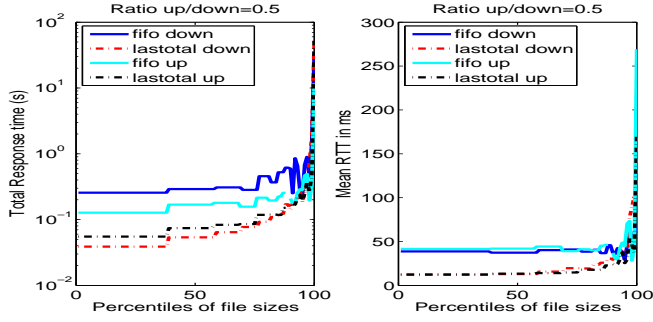


Figure 9: CoV=5.5, Input load=5 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 0.5$




Figure 10: CoV=5.5, Input load=5 Mbits/s, $\frac{\lambda_u}{\lambda_d} = 1$

queue, we have data flows whose service requirement follows a distribution G and ack flows whose service requirement follows the scaled distribution $\frac{S_{ack}}{2MSS}G$ since, on average, there is one TCP ack of size S_{ack} for every other MSS packet. We obtain the following formula for the conditional response time of the M/G/1/LASTOTAL queue, i.e., the response time of the downloads (see [18] for details):

$$T(x)_{LASTOTAL} = \frac{\hat{m}_2(x)}{2(1 - \hat{m}_1(x))^2} + \frac{x}{1 - \hat{m}_1(x)} \quad (4)$$

with

$$\hat{m}_1(x) = (\lambda_D + \frac{S_{ack}}{2MSS}\lambda_U)(m_1(x) + xF^c(x))$$

$$\hat{m}_2(x) = (\lambda_D + (\frac{S_{ack}}{2MSS})^2\lambda_U)(m_2(x) + x^2F^c(x))$$

We can observe from Equation (4) that the conditional response time under LASTOTAL is similar to the response time of a LAS system with only download traffic of intensity λ_D as in general $\lambda_D + \frac{S_{ack}}{2MSS}\lambda_U \simeq \lambda_D$ and $\lambda_D + (\frac{S_{ack}}{2MSS})^2\lambda_U \simeq \lambda_D$ (since S_{ack} is equal to 40 bytes, MSS is equal to 1500 bytes in general and $\lambda_U < \lambda_D$). Note however that the IP level capacity perceived by the downloads is affected by the upload traffic intensity.

As for the uploads, their data streams are indirectly controlled by the LASTOTAL scheduler when it schedules their TCP acks. The net result for those flows should be that their response time is close to the one of downloading flows, as seen in Section 6.

The main problem we face to check the validity of the

above queuing models is to correctly choose the IP level capacity of the access point. The next section will be devoted to this specific issue before we consider the issue of validating the models.

7.2 Wireless network capacity

We present in this section a technique used to infer the IP level bandwidth at the access point based on simulation results. Our technique to estimate the IP level capacity of the access point relies on the M/G/1/LASTOTAL response time formula. Let us consider a distribution with a minimum flow size k . Then, from Equation (4), we obtain that the response time of flows of size k is given by:

$$T(k)_{LASTOTAL} = \frac{\lambda(k/C)^2}{2(1 - \lambda k/C)^2} + \frac{k/C}{1 - \lambda k/C} \quad (5)$$

where $\lambda = \lambda_d + \frac{MSS}{2S_{ack}}\lambda_u$

The simple form of Equation (5) follows from the fact that $m(k) = \int_k^k u f(u) du = 0$, $m_2(k) = \int_k^k u^2 f(u) du = 0$ and $F^c(k) = 1$. Note that Equation (5) is identical for any distribution with a minimum flow size k . Using Equation (5), we can obtain an estimator \hat{C} of C by resolving the corresponding quadratic equation. A closed form formula for \hat{C} is given by Equation (6):

$$\hat{C} = \frac{\lambda k}{1 - \frac{1}{\sqrt{2\lambda T(k)_{LASTOTAL} + 1}}} \quad (6)$$

We plot in Figure 11 the estimated capacity values for all offered load, CoV and $\frac{\lambda_u}{\lambda_d}$ values. For a given distribution (a given CoV), the estimated capacity for $\frac{\lambda_u}{\lambda_d} = 0.5$ tends to be slightly larger than the one obtained for $\frac{\lambda_u}{\lambda_d} = 1$. This might be due to a higher level of contention on the wireless medium, as for $\frac{\lambda_u}{\lambda_d} = 1$, wireless stations have more traffic to send. We can also observe that the capacity values in Figure 11 are larger than the throughputs of the download traffic in Figures 3 and 4 (for a Pareto distribution with a CoV of 5.5). Note that one should consider only the bytes transmitted for the first time and subtract the retransmitted bytes as they are lost at the buffer of the access point.

One might argue that the reasoning made above with the M/G/1/LASTOTAL queue could be done with the M/G/1/PS queue. However, a key difference between LASTOTAL and FIFO in a WLAN is that short flows experience losses under FIFO while they don't under LASTOTAL in general. As a consequence, the response time of the short flows under FIFO is a function of both the IP level capacity of the access point and the losses they experience, which would lead to underestimate the IP level capacity. In contrast, under LASTOTAL, the response time of the short flows is almost exclusively a function of the IP level capacity.

7.3 Queuing Models for the Response time

We investigate here the accuracy of the queuing models proposed in Section 7.1. For different offered load values, we computed the response time at the access point queue per flow by subtracting to the one-way delay of each packet the latency and the transmission times on the wired section of the path.

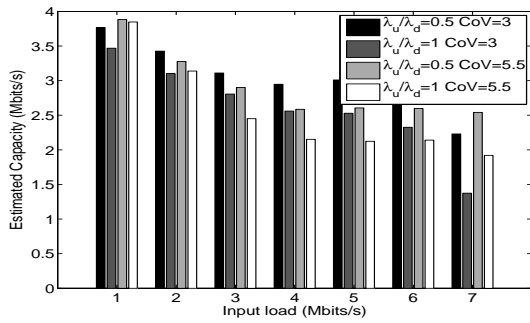


Figure 11: Estimated IP level capacities

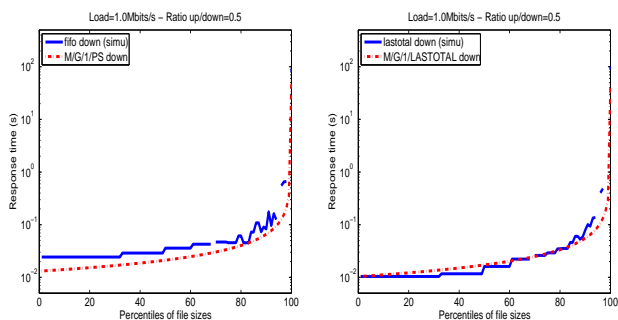
We compare in Figures 12 and 13 the theoretical models of Equations (4) and (2) with the results obtained by simulations for load values of 1 and 2 Mbits/s, a Pareto distribution with a CoV of 5.5 and a ratio $\frac{\lambda_u}{\lambda_d} = 0.5$. Overall, we observe that both models capture fairly well the response time at the access point. The M/G/1/PS appears less accurate when it comes to estimate the response time of the smaller flows. This is most probably because the model fails to capture the higher retransmission rates incurred by smaller flows under the FIFO policy.

For loads larger than 3 Mbits/s, the M/G/1/LASTOTAL model remains quite accurate as exemplified by Figure 14 for an input load of 5 Mbits/s. We have no result for the M/G/1/PS queue at this load, as the PS queue is in overload. Note that the M/G/1/LASTOTAL queue is also in overload. However, under LASTOTAL performance degrades gracefully when the load is above 1: these are first the largest flows that have an infinite response time while the smallest flows still have a finite response time³. Especially, the M/G/1/LASTOTAL model does not deliver a result for the last quantiles (above the 99th quantile) of the distribution as those largest flows experience an infinite response time in the queuing theory.

8. CONCLUSION

In this paper, we have investigated the use of a specific

³See [15] for details about LAS in overload - results can be easily extended to the case of LASTOTAL



(a) Response time at the access point (1 Mbits/s) for FIFO
(b) Response time at the access point (1 Mbits/s) for LASTOTAL

Figure 12:

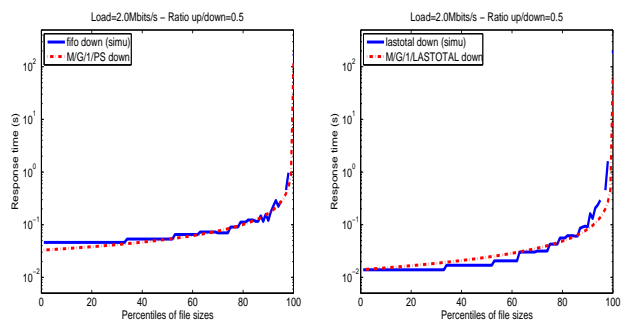
size-based scheduling policy, LAS, to both enforce fairness and provide small response times to most flows in a WLAN. We first underscore the need to simultaneously account for both directions of a connection due to the half duplex nature of 802.11 links. We call LASTOTAL the resulting implementation of LAS. LASTOTAL needs to be deployed at the access point only, without any modification of the MAC layer. We demonstrate using realistic synthetic workloads, that LASTOTAL solves the unfairness issue due to DCF in 802.11 WLANs and ensures small response times to the majority of the flows under any load conditions.

We further introduce and evaluate single server queuing models of the response time at the access point for both FIFO and LASTOTAL. A key problem when assessing the validity of those models is to correctly set the IP level capacity of the access point. We presented a method to estimate the capacity from the simulations. The results we have obtained are encouraging as queuing models turn out to be accurate under various load conditions.

In terms of future work, we would like to extend our analytical model of the response time at the access point to the global response time of connections (from the first to the last packet of a connection). This is a complex problem as one needs to account for the waiting time experienced by TCP acks on the reverse path of the connection, which is difficult to evaluate due to the shared nature of the 802.11 link. We would also like to investigate the benefits of LASTOTAL in ad-hoc and wireless mesh networks where the performance of TCP flows are subject to wide variations in general.

9. REFERENCES

- [1] BIANCHI, G. Performance analysis of the ieee 802.11 distributed coordination function. *Selected Areas in Communications, IEEE Journal on* 18, 3 (2000), 535–547.
- [2] BOTTIGLELIENGO, M., CASETTI, C., CHIASSERINI, C., AND MEO, M. Short-term fairness for tcp flows in 802.11b w lans. In *INFOCOM 2004* (2004), vol. 2, pp. 1383–1392 vol.2.
- [3] BRUNO, R., CONTI, M., AND GREGORI, E. Analytical modeling of tcp clients in wi-fi hot spot networks. In *NETWORKING* (2004), pp. 626–637.
- [4] BRUNO, R., CONTI, M., AND GREGORI, E. Performance modelling and measurements of tcp



(a) Response time at the access point (2 Mbits/s) for FIFO
(b) Response time at the access point (2 Mbits/s) for LASTOTAL

Figure 13:

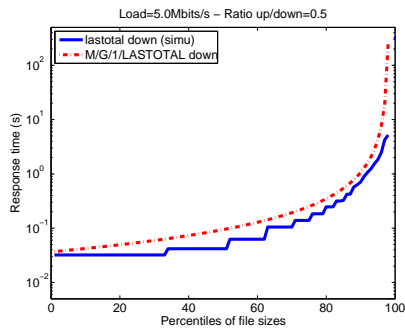


Figure 14: Response time at the access point (5 Mbits/s) for LASTOTAL

transfer throughput in 802.11-based wlan. In *MSWIM '06* (2006), pp. 4–11.

- [5] CROVELLA, M. E., ET AL. *A Practical Guide to Heavy Tails*. Chapman and Hall, New-York, 1998, ch. 3.
- [6] HARCHOL-BALTER, M., SCHROEDER, B., BANSAL, N., AND AGRAWAL, M. Size-based scheduling to improve web performance. *ACM Trans. Comput. Syst.* *21*, 2 (2003), 207–233.
- [7] HERNANDEZ-CAMPOS, F., KARALIOPOULOS, M., PAPADOPOULI, M., AND SHEN, H. Spatio-temporal modeling of traffic workload in a campus wlan. In *WICON '06* (2006).
- [8] HEUSSE, M., URVOY KELLER, G., AND DUDA, A. An experimental comparison of layer 2 and 3 mechanisms for improving user perceived performance of 802.11 wireless lans. Tech. Rep. EURECOM+220, Eurecom, France, May 2008.

- [9] HU, M., ZHANG, J., AND SADOWSKY, J. Size-aided opportunistic scheduling in wireless networks. In *GLOBECOM '03. IEEE* (2003).
- [10] KLEINROCK, L. *Queuing Systems, Volume II: Computer Applications*. Wiley, New York, 1976.
- [11] LOPEZ-AGUILERA, E., HEUSSE, M., GRUNENBERGER, Y., ROUSSEAU, F., DUDA, A., AND CASADEMONT, J. An Asymmetric Access Point for Solving the Unfairness Problem in WLANs. *IEEE Transactions on Mobile Computing* (2008). Preprint.
- [12] MENG, X. G., WONG, S. H. Y., YUAN, Y., AND LU, S. Characterizing flows in large wireless data networks. In *MOBICOM'04* (2004), pp. 174–186.
- [13] NUYENS, M., AND WIERMAN, A. The foreground-background queue: A survey. *Perform. Eval.* *65*, 3-4 (2008), 286–307.
- [14] RAI, I. A., URVOY-KELLER, AND BIRSACK, E. W. Las scheduling approach to avoid bandwidth hogging in heterogeneous tcp networks. In *HSNMC'04* (July 2004).
- [15] RAI, I. A., URVOY-KELLER, G., AND BIRSACK, E. W. Analysis of las scheduling for job size distributions with high variance. In *Proc. ACM SIGMETRICS* (June 2003), pp. 218–228.
- [16] SCHRAGE, L. E. The queue m/g/1 with feedback to lower priority queues. *Management Science* *13*, 7 (1967), 466–474.
- [17] SINHA, P., PILOSOFF, S., RAMJEE, R., RAZ, D., AND SHAVITT, Y. Understanding TCP fairness over wireless LAN. In *INFOCOM 2003* (Apr. 2003).
- [18] URVOY KELLER, G., AND BEYLOT, A.-L. Improving flow level fairness and interactivity in WLANs using size-based scheduling policies. Tech. Rep. EURECOM+2373, Eurecom, France, Nov 2007.