

# A Flexible and Green Scheduler for providing QoS and High Throughput over Wireless Links

Maurizio Casoni<sup>1</sup>, Carlo Augusto Grazia<sup>1</sup>, Paolo Valente<sup>2</sup>

<sup>1</sup>Department of Engineering *Enzo Ferrari*, <sup>2</sup>Department of Physics, Computer Science and Mathematics  
University of Modena and Reggio Emilia, via Vignolese 905, 41125 Modena,  
{maurizio.casoni, carloaugusto.grazia, paolo.valente}@unimore.it

**Abstract:** Realizing a communication system over wireless links that simultaneously offers QoS guarantees, high throughput and low energy consumption is a difficult task. A common solution is using a single, *integrated* scheduler that deals both with the QoS guarantees and the wireless link issues. Unfortunately, such an approach is little flexible and does not allow any of the existing high-quality schedulers for wired links to be used without modifications. To address these issues, in this paper we validate a modular architecture which permits the use, as they are, of existing packet schedulers for wired links over wireless links, and at the same time allows the flexibility to adapt to different channel conditions.

We validate the effectiveness of the modular architecture by showing, through experimental results, that this architecture enables us to get a new scheduler with the following features, just by combining existing schedulers: execution time and energy consumption close to that of just a Deficit Round Robin, accurate fairness and delay guarantees, possibility to set, by changing one parameter, the desired trade-off between throughput-boosting level and granularity of the service guarantees. In particular, we show that this scheduler, which we named High-throughput Twin Fair scheduler (HFS), outperforms one of the most accurate and efficient integrated schedulers available in the literature

**Index Terms**—cross-layering, energy consumption, QoS, scheduling, throughput

## I. INTRODUCTION

State-of-the-art wireless technologies aim at supporting increasingly high data rates for applications such as video streaming, web browsing and file sharing, in both stationary and nomadic/mobile scenarios. Moreover, given the growing spread of smartphones and tablets, an increasing number of users access to wireless networks everyday.

This trend puts several limits on how network and/or service providers can effectively supply adequate quality of service (QoS) to their users, since over-provisioning is hard to implement in wireless contexts. Most current wireless systems directly provide QoS capabilities, in terms of traffic differentiation, traffic prioritization and so on. In this respect, one of the most important sub-systems involved in QoS provision is the packet scheduler, which properly sets the order by which packets are sent over a given interface, both in the uplink and in the downlink. State-of-the-art solutions to provide both QoS guarantees, as well as a high throughput, are based on *cross-layering*: scheduling decisions are made by using also channel state information coming from the MAC layer [1], [2], [3], [4]. For example, per-destination channel conditions

*This research was funded by the Italian Ministry of Education University and Research through the PRIN 2009 project SFINGI (Software router to Improve Next-Generation Internet).*

DOI: 02.WCMCS.2013.1.34

© Association of Computer Electronics and Electrical Engineers, 2013

may be considered when deciding which flow to serve, in order to avoid transmission failures. In these proposals, just *one integrated* scheduler takes *all* the scheduling decisions, based on these issues and on the desired QoS.

Unfortunately, integrated solution entails a few drawback. High-quality schedulers for reliable links [5], [6], [7], [8], [9], [10], [11], [12] cannot be used and converted into cross-layer schedulers, without modifying them. After these modifications, the guarantees provided by the scheduler are likely to change and need to be recomputed. Finally, if the medium access protocol or the channel technology changes, or if we want to use new techniques for achieving an even higher throughput or saving energy, then the scheduler is likely not to fit the new technology or requirements. Hence it may need to be modified again. Especially, if we want to use the same scheduler on heterogeneous wireless technologies, we may need to define a different version of it for each technology.

In this work we validate the flexible and effective packet-scheduling architecture proposed in [13]. The architecture focuses on *local* packet schedulers, which are executed inside wireless nodes, and decide (only) the order by which packets are transmitted over the outgoing links of the node. Such architecture preserves both effectiveness and flexibility, and permits to reuse existing schedulers without modification; we implement such architecture in which the scheduler just chooses the next packet to transmit according to its QoS policy, but delivers it to a *middle layer* instead of the MAC layer. The middle layer then deals with the issues of the wireless medium and reorders packet transmissions if needed. With this architecture it is easy to cherry pick from the literature and combine the best solutions in terms of service guarantees, computational cost, power consumption and throughput boosting. As for the last two goals, we note that this architecture allows a system to *profit from cross-layering* while still *preserving flexibility*.

#### A. Contribution

How effective is this architecture? In this paper we answer, in a sense, to this question by defining a new flexible, efficient and green packet scheduler for wireless links: High-throughput twin Fair Scheduler (HFS). To present it more thoroughly we complete the architecture description with the developed mechanism that moves the packets from the QoS layer to the middle layer. This mechanism is implemented in a software module called packet prefetcher.

HFS provides two contributions to energy reduction: first, by increasing the throughput, it increases the number of packets successfully transmitted per energy consumed (the number of retransmission is also lower), second, according to our experiments, the time and energy needed to execute HFS for each packet enqueue/dequeue are close to those of just DRR. Moreover, while still reducing energy consumption and boosting the throughput, we show also that HFS preserves high QoS guarantees.

#### B. Organization of the document

In Section II we describe the architecture proposed in [13]. In Section III we show the testbed deployed to validate the efficiency of the architecture. Then we show HFS, a new packet scheduler for wireless links based on such architecture and implemented in such testbed in Section IV. In Section V, we validate HFS performance by comparing it with the best high-performance schedulers for wired links and with the best integrated scheduler for wireless links, showing experimental results. Finally, in Section VI we highlight our conclusions.

## II. ARCHITECTURE

The architecture proposed in [13] is shown in Figure 1 for a system containing only one outgoing link. The architecture is made of two layers: a *QoS Provisioning Layer*, hereafter called just QoS layer for short, and a *MAC Scheduling & Abstraction Layer*, hereafter abbreviated as MAC-SAL. The purpose of the first layer is putting together, conceptually, the two most important components for providing service guarantees over a transmission link: the packet classifier and the packet scheduler. Packets are passed from the QoS layer to the MAC-SAL by a *packet prefetcher*, described in detail in this section. For a detailed description of both the QoS layer and the MAC-SAL layer we remind to [13].

*Packet prefetcher.* The higher is the number of non-empty MAC (flow) queues, the higher is the number of head packets among which the MAC scheduler can choose the next packet to transmit. Hence, the higher is the probability that the MAC scheduler can pick good packets with respect to the goals it wants to achieve. For example, suppose that the goal of the MAC scheduler is to keep a high throughput and that some MAC queues contain packets for destinations with bad channel conditions. If many/few other queues are not empty,

then the probability for the scheduler to have at its disposal better packets to transmit is high/low. In the end, to maximize the effectiveness of the MAC-SAL, its queues should be always kept as full as possible. This is exactly the purpose of the *packet prefetcher* depicted in the middle of Figure 1.

The behavior of the prefetcher depends on a parameter  $Q$ , measured in number of bytes. We say that the prefetcher prefetches a packet if it invokes the function *get\_next\_pkt()* to get the next packet, say  $pkt$ , from the QoS layer and then invokes *MAC-SAL-send(pkt)* to insert the packet in the MAC-SAL. We assume that a prefetched packet is never dropped by the MAC-SAL. Finally, we denote as  $S(t)$  the sum of the sizes of the packets queued in the MAC-SAL at time  $t$ .

The prefetcher operates on each of the following two events: when a new packet arrives in the QoS layer and when a new packet is dequeued from the MAC-SAL. On the first event, let  $pkt$  be the new arriving packet. If  $S(t) < Q$  when  $pkt$  arrives, then the packet prefetcher immediately prefetches  $pkt$ . On the other event, if  $S(t)$  becomes lower than  $Q$  when the packet is dequeued from the MAC-SAL, then the packet prefetcher starts prefetching new packets until  $S(t) < Q$  does not hold any more. In other words, the combination of the set of queues in the MAC-SAL and of the packet prefetcher implements a *shared buffer with virtual queueing*, a device in which the memory used to store the packets is shared and the queues are only virtually separated.

We conclude by noting that this scheme is much more effective than an approach with distinct queues and with the same total memory. As for the second scheme, suppose for example that all the queues have the same length, equal to  $Q/M$  plus one maximum packet size  $L$ . If all the packets prefetched in a given time interval are destined to the same MAC queue, then the queue becomes full only after  $Q/M + L$  bytes have been prefetched. If the next packet to prefetch is again for the same queue, then the prefetcher must block. In the shared-buffer scheme the prefetcher must block only after  $Q + L$  bytes have been prefetched, which increases the probability that more queues are not empty.

### III. TESTBED

In this section we present the test environment used to easily deploy, test, modify and validate flexible and efficient packet schedulers for wireless links as well as scenarios to model the radio channel characteristics.

We ran our experiments using the test environment [14], which is an improved version of the original test environment [15], extended to test schedulers over the modular architecture solution showed in Section II and to simulate also wireless link characteristics. With the help of this flexible tool it is possible to measure and compare one against each other the real execution time, simulated throughput and QoS performance of different schedulers. An existing packet scheduler can be easily plugged into this environment, after at most some little interface changes. Inside the environment, the scheduler can then be exercised with the desired sequence of enqueue/dequeue requests, through a *controller* that iteratively switches between two phases: an enqueue phase in which it generates fake packets by picking them from a free list, and a dequeue phase in which it dequeues packets from the scheduler and reinserts them into the free list. The switch occurs according to two configurable max-total-backlog and min-total-backlog thresholds.

*Configuration*:- Each test consisted of  $10M^1$  events with a typical balancing of 5M packet enqueues and 5M packet dequeues, with the controller configured so as to let flows oscillate between null backlog and a backlog of 10 packets each. Such an enqueue/dequeue pattern happened to be the most demanding one for the schedulers. Packets had a fixed size of about 1700 bytes, with no cache-line alignment. The payload of the packets was never either read or written. No migration and no packet drop occurred in any run.

*Arrival Pattern*:- The controller switches between enqueue and dequeue mode to control the number of pending packets of each run. Packets arrival pattern is set in a *bursty* manner:

the packets generated by the controller are sent to the flows with bursts that have random size proportional to the flow weight. Also a single threshold per-flow has been setted, if a flow drops below a certain threshold of filling, the controller refill it with a given probability

*Wireless scenarios*. This tool can be executed to simulate a wireless link. To run our tests over the modular architecture we considered a specific scenario described here with its specific well known wireless parameters simulated, that is:

- Total bandwidth: 54 Mb/s;
- 20 Subscriber Stations (SS);
- 2.7 Mb/s per SS;
- For each SS, 5 flows:

---

<sup>1</sup>The package of the test environment [14] contains the script, run test.sh, that we used to run the tests.

- One with ~ 1.6 Mb/s reserved for video or VoIP (20 flows with weight 20);
- Another with ~ 0.8 Mb/s reserved for WEB browsing or direct download (20 flows with weight 10);
- The other three with ~ 0.1 Mb/s reserved for download/sharing (3·20 flows with weight 1).

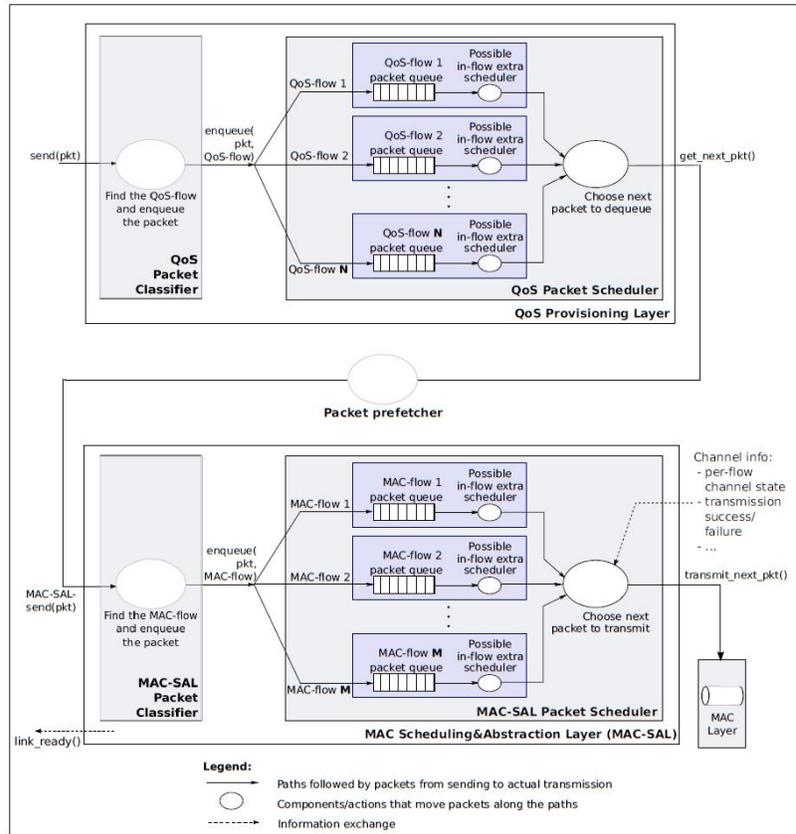


Fig. 1. Modular architecture of providing QoS over a wireless link

As for base stations channel condition, we set for each of them a packet loss probability ( $P_{loss}$ ) to emulate different user conditions. In fact, we let  $P_{loss}$  ranging linearly from  $10^0$  to  $10^{-1}$  and we considered also  $P_{loss}$  of  $10^{-2}$ ,  $10^{-3}$  and  $10^{-4}$ . We considered static channel condition in our tests to evaluate possible phenomena for different user conditions. From an user classification point of view, we define a threshold to distinguish user in *good* and *bad* channel condition. This threshold is placed at 20% of packet loss. Users/flows with less or equal to 20% of packets lost are considered as users/flows in good conditions, while users/flows with more than 20% of packets lost are considered as user/flows in bad conditions. This because above 20% most applications do not work properly; for instance, both the TCP window and the VoIP controller mechanism do not perform well under the aforementioned condition.

**Statistics:-** To measure, simulate and validate schedulers' performance we implemented different well-known metrics. We used all of these indicators to easily design, test and validate different packet scheduler solutions. Here is a list of the main performance metrics and a detailed description:

**Throughput:-** to measure the throughput we simulated the *normalized throughput* achieved by the schedulers. This parameter is the amount of successfully transmitted packets for each flow divided by the amount of total sent packets. We computed the normalized throughput as

$$thr \equiv \frac{\sum_i pkt_{sent_i} (1 - P_{loss_i})}{\sum_i pkt_{sent_i}} \quad \text{where } pkt_{sent_i} \text{ is the number of packets sent by}$$

the flow  $i$ ,  $P_{loss_i}$  is the packet loss probability of the flow  $i$  and  $pkt_{sent_i} (1 - P_{loss_i})$  is the number of successfully transmitted packets by the flow  $i$ ;

*Execution time*:- how to measure the execution time? At the end of each run, we measured the total execution time of the run and divided it by the total number of enqueues, or equivalently of dequeues executed. We obtained therefore the average total packet-processing time, i.e., the execution time of an enqueue plus a dequeue, inclusive also of the cost of generating and discarding an empty, fixed-size packet;

*Energy consumption*:- according to the models in [16], [17], lower/higher relative execution times imply also lower/higher relative energy consumptions. Moreover, the increase of the throughput imply the increase of the number of packets successfully transmitted per energy consumed. In this sense we will consider the energy consumption level according to the execution time and the throughput level;

*Time Worst-case Fair Index (T-WFI)*:- another useful QoS guarantees metric is the T-WFI, which allows to evaluate, in a single value, both fairness and delay: on one hand it shows the fairness, in terms of delay from the worst-case completion time of a packet in a perfectly fair system, while on the other hand it allows to instantly calculate the actual delays incurred by packets depending on the occupation of queues. In a perfectly fair system, the worst-case completion time of a packet is equal to its queue length (including the packet itself) divided by the packet's flow guaranteed ratio. Assuming that the link rate  $R$  is constant, we

measured T-WFI $_k$  for flow  $k$  as:  $T - WFI \equiv \max \left( t_{deq} - t_{enq} - \frac{Q^k(t_{enq})}{\phi^k K} \right)$  where  $t_{deq} - t_{enq}$  is the delay

experienced by a packet  $pkt$ ,  $t_{enq}$  is the number of packets dequeued by the system when  $pkt$  is enqueued, while  $t_{deq}$  is the number of packets dequeued by the system when  $pkt$  is dequeued.  $Q^k(t_{enq})$  is the backlog of flow  $k$  just after the arrival of the packet and  $\phi^k$  is the relative weight of the flow.

*Test equipment*:- We ran our tests on two systems with the following software and hardware characteristics:

- Ubuntu 12.04.2, 64-bit kernel 3.2.0, Intel Core Dual- E2200 @ 2.20GHz, gcc 4.6.3 -O3;
- OS X 10.7.5, Darwin 11.4.0, Intel Core i5-2557M @ 1.8 GHz, gcc 4.2.1 -O3.

Since the relative performance of the schedulers were about the same on the two systems, we report our results only for the first system. In the next two sections we first show how we used the test environment described here to define HFS packet scheduler, and then we show how we validated its performance.

#### IV. HFS

Once described the architecture, it is easy to describe HFS. We considered the best schedulers for wired links available in the literature, that is:

- WF<sup>2</sup>Q+: an optimal service guarantees with  $O(\log n)$  complexity [5];
- DRR: a scheduler with extremely low time complexity  $O(1)$ , but with  $O(n)$  deviation form optimal service [7];
- QFQ+: a quasi-optimal service guarantees scheduler with execution time close to the DRR one [18].

To develop HFS we combined these schedulers placing them at QoS layer and/or at MAC-SAL layer. For a matter of space we do not show all the results of these combinations (we have left a detailed analysis of all solutions in a different work), but we present directly our solution with QFQ+ both at QoS and MAC-SAL layer. QFQ+ achieves quasi-optimal service guarantees at QoS level while still preserving low execution time overhead, on the other side QFQ+ placed at MAC-SAL layer achieves high-throughput with, also in this case, a low time complexity. In brief, to understand how HFS is composed by, just look at the Figure 1 and place QFQ+ at the end of both layers inside the packet scheduler box. We validated HFS in

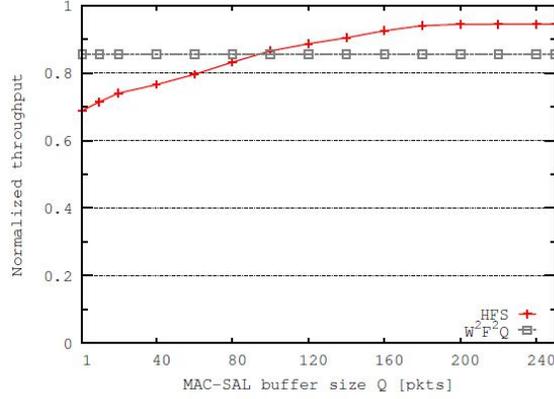


Fig 2. Normalized throughput for HFT and  $W^2F^2Q$

Section V comparing its performance with the best existing schedulers for wired networks and with the best integrated one for wireless links, in particular we compared its execution time against DRR, its throughput level against  $W^2F^2Q$  and its service guarantees against  $WF^2Q+$ .

## V. RESULTS

Here we validate the efficiency of HFS packet scheduler by comparing its performance with the best packet schedulers for wired links and with the best integrated scheduler for wireless links. To validate HFS results we considered different schedulers as a benchmark:

- $W^2F^2Q$ , to validate HFS throughput results against the best integrated scheduler available in the literature [19];
- *DRR*, to validate HFS execution time against the best high-performance scheduler for wired links in terms of time complexity;
- $WF^2Q+$ , to validate HFS guarantees against the best high-performance scheduler for wired links in terms of service guarantees.

Let us call *double-SCHED* a double instance of the scheduler *SCHED* compliant with the modular architecture with *SCHED* placed both at QoS and at MAC-SAL layer (e.g. double-DRR is obtained placing DRR scheduler both at QoS level and at MAC-SAL level)<sup>2</sup>. Similarly, we will call *single-SCHED* a single instance of the scheduler *SCHED* compliant with the classical architecture.

*Throughput.* The integrated scheduler  $W^2F^2Q$  models temporary bursty channel errors of a wireless link only, with a two state Markov chain, to simulate flows in good or bad channel conditions. Unfortunately, such a distinction does not hold in our model, since we considered a scenario in which a flow can experience long term static channel conditions as well, based for instance on the position of the user. Anyway, comparing  $W^2F^2Q$  with HFS can benefit the first in terms of throughput, since users in good state manage the total bandwidth available leaving flows in bad state with null weight. We simulate this kind of event by using a threshold in our environment, where flows with more than 20% of packet loss are classified as flows in bad state with weight equal to zero, while the other flows are classified as flows in good state with weight equal to their own weight according to the weight distribution policy. Figure 2 shows how in this favourable case  $W^2F^2Q$  scheduler obtains high normalized throughput.

However, HFS achieves better throughput performance with respect to  $W^2F^2Q$  for  $Q \geq 100$  due to his fine-grained choice among good flows, which increases with the MAC-SAL shared buffer size  $Q$ .

*Execution time.* To evaluate HFS in terms of execution time (and energy consumption) we use as benchmark the scheduler double-DRR, which is the simplest possible solution to achieve high throughput preserving QoS guarantees with a low computational cost of  $O(1)$  with our modular architecture. Figure 3 shows that HFS has a really close execution time to the double instance of DRR which is the best high-performance scheduler for wired links in terms of execution time. Furthermore, the picture shows that the packet-processing time of HFS is compliant with the link rate (i.e. much lower than the packet transmission time). Why we do not considered  $W^2F^2Q$  in the execution time graph?  $W^2F^2Q$  scheduling complexity is  $O(n)$ , too high for schedulers in backbone network where  $n$  is very large and more than the HFS scheduling complexity which is close to the DRR optimal cost of  $O(1)$ .

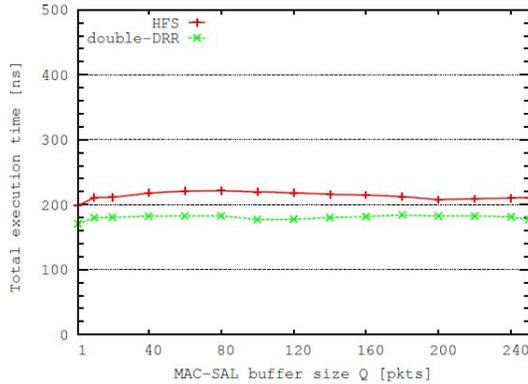


Fig 3. Execution time of HFS against double DRR

*Energy consumption.* According to the models in [16], [17], lower/higher relative execution times imply also lower/higher relative energy consumptions. It is the case of HFS, as showed in Figure 3, with a really close execution time to DRR which is the best high-performance packet scheduler for wired links in terms of execution time. Moreover, in Figure 2 we show that HFS achieves also higher throughput then the best integrated packet scheduler for wireless links, which is  $W^2F^2Q$ . In this way, by increasing the throughput, HFS increases the number of packets successfully transmitted per energy consumed (the number of retransmission is also lower). Therefore, reducing the execution time and boosting the throughput permits to HFS to reduce the energy consumption.

*QoS guarantees.* As regards QoS metric, we show only Time Worst-case Fair Index because it allows to evaluate, in a single graph, both fairness and delay: on one hand it shows the fairness, in terms of delay from the worst-case completion time of a packet in a perfectly fair system, while on the other hand it allows to instantly calculate the actual delays incurred by packets depending on the occupation of queues. In a

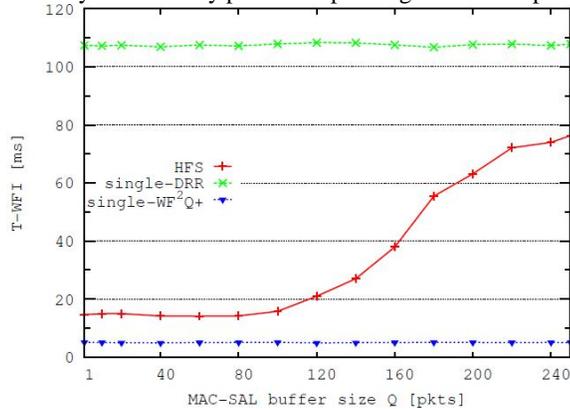


Fig 4. T-WFI for HFS, WF2Q+ and DRR for Flows in intermediate conditions

perfectly fair system, the worst case completion time of a packet is equal to its queue length (including the packet itself) divided by the packet's flow guaranteed ratio Figure 4 validates the performance of HFS in terms of QoS for flows in *intermediate conditions*, i.e. for flows/users with at least 0.8 of normalized throughput per-flow (HFS guarantees this kind of worst case performance if the station loses less or equal to 20% of packets). With a packet loss ratio above 20% most applications do not work properly; for instance, both the TCP window and the VoIP controller mechanism do not perform well under the aforementioned condition. Why we do not considered  $W^2F^2Q$  in QoS-guarantees graph? As we have said, the algorithm  $W^2F^2Q$  model temporary bursty channel errors of a wireless link only. This way the scheduler tries to implement the *compensation* mechanism to guarantees long term fairness between different flow-condition services under the hypothesis that an error-prone flow has sufficient time to make up for their lag after recovery of channel (see [19]). Unfortunately, such assumption does not hold in our model because we

considered a scenario in which a flow can experience also long term static channel condition, based on the position of the user. In this sense, comparing  $W^2F^2Q$  with our model from a QoS guarantees point of view can be unfair because users/flows classified in bad state can lag forever with quasi-zero service guarantees. By comparing Figure 2 and 4 we can say that it is possible to choose the desired trade-off between throughput-boosting level and granularity of the service guarantees, by only setting the parameter  $Q$ . For instance, by setting a value of  $Q$  equal to 100 packets, HFS reaches a normalized throughput close to 90%, greater than  $W^2F^2Q$  one, while still preserving service guarantees close to the optimal ones of  $WF^2Q$ .

## VI. CONCLUSIONS

In this paper we have validated a general, modular architecture for decoupling the task of providing QoS guarantees from the task of dealing with the idiosyncrasies of a wireless link. We defined a new test-environment which easily permits to test existing packet schedulers and analyze the result in terms of execution time, QoS guarantees and also throughput achieved over wireless links. We also define HFS, a new flexible, efficient and green packet scheduler for providing low energy consumption, high throughput and QoS guarantees over wireless links. As a validation of HFS we show, through experimental results, the high performance of this new packet scheduler with execution time and energy consumption close to that of just a Deficit Round Robin, higher throughput than the best integrated scheduler Wireless Worst-case Fair Weighted Fair Queuing and accurate fairness and delay guarantees close to the optimal ones of Worst-case Weighted Fair Queuing.

## REFERENCES

- [1] T. S. E. Ng, I. Stoica, and H. Zhang, "Packet fair queuing algorithms for wireless networks with location-dependent errors," *Proceedings of IEEE INFOCOMM 98*, vol. 3, no. c, pp. 1103–1111, 1998.
- [2] S. Lu, V. Bharghavan, and R. Srikant, "Fair scheduling in wireless packet networks," *IEEE/ACM Trans. Netw.*, vol. 7, no. 4, pp. 473–489, 1999.
- [3] Y. Yi, Y. Seok, T. Kwon, Y. Choi, and J. Park, "W2f2q: packet fair queuing in wireless packet networks," in *Proceedings of the 3rd ACM international workshop on Wireless mobile multimedia*, ser. WOWMOM '00. New York, NY, USA: ACM, 2000, pp. 2–10.
- [4] A. Iera, A. Molinaro, and S. Pizzi, "Channel-aware scheduling for qos and fairness provisioning in ieee 802.16/wimax broadband wireless access systems," *Ieee Network*, vol. 21, no. 5, pp. 34–41, 2007.
- [5] H. Bennet, Jon C. R. and Zhang, "Hierarchical packet fair queuing algorithms," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 675–689, 1997.
- [6] P. Valente, "Exact gps simulation and optimal fair scheduling with log-arithmetic complexity," *IEEE/ACM Transactions on Networking*, vol. 15, no. 6, pp. 1454–1466, 2007.
- [7] M. Shreedhar and G. Varghese, "Efficient fair queuing using deficit round robin," in *IEEE/ACM Transactions on Networking*, 1995, pp. 375–385.
- [8] C. Guo, "SRR: An  $O(1)$  time complexity packet scheduler for flows in multi-service packet networks," in *In ACM SIGCOMM 2001*, 2001, pp. 211–222.
- [9] S. Ramabhadran and J. Pasquale, "The stratified round robin scheduler: design, analysis and implementation," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1362–1373, 2006.
- [10] X. Yuan and Z. Duan, "Fair round-robin: A low complexity packet scheduler with proportional and worst-case fairness," *IEEE Transactions on Computers*, vol. 58, no. 3, 2009.
- [11] D. C. Stephens, J. C. Bennett, and H. Zhang, "Implementing scheduling algorithms in high-speed networks," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 6, pp. 1145–1158, Jun 1999.
- [12] M. Karsten, "Approximation of generalized processor sharing with stratified interleaved timer wheels," *IEEE/ACM Transactions on Networking*, vol. 18, no. 3, pp. 708–721, June 2010.
- [13] M. Casoni, A. Paganelli, and P. Valente, "A modular architecture for qos provisioning over wireless links," in *Proc. of the 8th IEEE International Workshop PAEWN*, Barcelona, (Spain), 2013.
- [14] <http://algroup.unimore.it/people/paolo/agg-sched/test-env.tgz>.
- [15] [http://info.iet.unipi.it/\\_luigi/papers/20100210-qfq-test.tgz](http://info.iet.unipi.it/_luigi/papers/20100210-qfq-test.tgz).
- [16] A. Bartolini, M. Cacciari, A. Tilli, and L. Benini, "A distributed and self-calibrating model-predictive controller for energy and thermal management of high-performance multicores," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, 2011, pp. 1–6.
- [17] M. Sadri, A. Bartolini, and L. Benini, "Single-chip cloud computer thermal model," in *Thermal Investigations of ICs and Systems (THER-MINIC), 2011 17th International Workshop on*, 2011, pp. 1–6.
- [18] P. Valente, "Providing near-optimal fair-queuing guarantees at round-robin amortized cost," to appear in the *Proc. of the 22nd International Conference on Computer Communications and Networks, ICCCN 2013*.
- [19] Y. Yi, Y. Seok, T. Kwon, Y. Choi, and J. Park, "W2f2q: packet fair queuing in wireless packet networks," in *WOWMOM*, 2000, pp. 2–10.