Compensation Modeling for QoS Support on a Wireless Network

Stefan Bucheli Jay R. Moorman John W. Lockwood Sung-Mo Kang

Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

Abstract

This paper presents the design of a new error compensation model for providing QoS support on a wireless network. The compensation model uses different compensation strategies for the different priority classes in conjunction with a Multiclass Priority Fair Queuing (MPFQ) algorithm. Criteria for fair error compensation and a classification of error compensation models are established in order to show the properties of different models, and to select the best error compensation for each priority class. Simulation results show that the new MPFQ compensation model meets the long-term fairness guarantees and provides an improved flow separation.

1 Introduction

As wireless networking grows, the issue of providing fair channel access among multiple hosts over a shared wireless channel is becoming more and more important. Limited bandwidth is shared among many flows, each with different requirements for bandwidth and delay. Unlike scheduling in wired networks, wireless packet scheduling must take channel errors into account since they are much more frequent on a wireless channel. Therefore, wireless scheduling algorithms use error compensation models to support and guarantee fair service.

The Multiclass Priority Fair Queuing (MPFQ) algorithm in [7] schedules packets according to Asynchronous Transfer Mode (ATM) service categories. Traffic flows in these categories have specific parameters for delay and bandwidth. These parameters are guaranteed throughout the lifetime of the connection by the network. The MPFQ scheduler uses these traffic parameters to extend the Quality of Service (QoS) into the wireless domain. Located in the wireless base station of a hybrid wired/wireless network, the scheduler provides the last link for end-to-end Quality of Service guarantees.

In this paper, we propose an error compensation model [3] for the MPFQ algorithm that compensates individual traffic categories differently. The MPFQ algorithm, with this error compensation model, has been implemented in the *ns* simulator. Simulations are presented to verify the improvements of this error compensation model.

2 Multiclass Priority Fair Queuing

2.1 Model

The network model considered in this paper consists of a high-speed wired ATM backbone that is extended to a packetized wireless cellular network. The wireless cellular network is divided into partially overlapping cells that transmit on different logical channels. Each cell contains a base station that is connected to the wired network. The base station schedules packets for transmission in its particular cell for both the uplink and the downlink channel. The channel is shared by all the mobile hosts and the base station within one cell, and its capacity is assumed to be constant.

The channel is said to be in error for a particular flow if the base station cannot communicate with the particular mobile using that flow. Errors are assumed to be location dependent, since one mobile host experiencing a fade will not affect the transmission of another mobile host at a different location. Errors are also assumed to be typically bursty in the wireless environment.

Priority	Traffic Type	Parameters	Queuing Policy
1	CBR	PCR, CTD	Weighted Fair Queuing
2	rt-VBR	SCR, CTD, MBS	Weighted Fair Queuing
3	nrt-VBR, ABR-MCR	SCR, MBS, MCR	WRR Scheduler
4	ABR	Unspecified	Recirculating FIFO
5	UBR	Unspecified	FIFO

Table 1: ATM Prioritized Traffic in MPFQ

2.2 MPFQ Priority Classes

The MPFQ algorithm schedules traffic flows in prioritized classes according to each ATM service category [1]. Each of the priority levels has its own packet scheduler that determines the packet order in its class. The algorithm uses a modified Weighted Fair Queuing (WFQ) policy in the higher priorities, Weighted Round Robin (WRR) scheduling for the middle priorities, and FIFO scheduling at the lower priorities (Table 1).

The packet to be transmitted is chosen from the highest priority class that has a packet available. This means that packets from a flow in a lower priority class must wait until there are no packets available to send in a higher class. With this prioritized approach, the traffic streams requiring short delays (CBR and rt-VBR) can meet their delay guarantees, while all the classes receive their bandwidth guarantees.

3 Error Compensation in Wireless Networks

3.1 Criteria for Evaluating Compensation Models

The goal of wireless scheduling is to provide fair guaranteed service to all traffic flows. However, it is not evident how fairness should be defined in a wireless network where channel errors may occur. To compare error compensation models, a set of criteria for wireless fair scheduling and error compensation are needed.

In recent publications, criteria to characterize fairness for wireless scheduling algorithms have been defined (*wireless fair service* *model* [9] and *Channel-condition Independent Fair (CIF)* [10]). These criteria include long-term and short-term throughput bounds, fairness among backlogged flows, and delay bounds. Additionally in [9], criteria to support delay-sensitive and error-sensitive flows, and to optimize the schedulable region were included. In order to find the appropriate error compensation model for the MPFQ algorithm, the following criteria are used in this paper:

- *Short-term fairness* among backlogged flows that perceive no channel error. These flows should receive their fair share of service in an arbitrary time interval, according to their weight.
- *Long-term fairness* among backlogged flows with bounded channel error. Over a longer time interval, all flows (except those with permanent errors) should receive their fair amount of service.
- *Delay bounds* for packets of any flow. These bounds depend on the nature of the location-dependent channel errors.
- *Short-term throughput bounds* for error-free flows. An error-free flow has to receive a minimum amount of service (rate), even if it has received excess service in the past.
- *Long-term throughput bounds* for all flows with bounded channel error.

3.2 Classification of Compensation Models

In this section, the concepts of error compensation models are classified according to typical characteristics. This classification is derived from the generic architecture of wireless scheduling algorithms described in [9]. Most compensation models in the existing wireless scheduling algorithms [4, 5, 10, 12] fit into this scheme.

A traffic flow is called *leading* if it has received more service than it would have received with no channel errors present in the system. A flow that has received less service than it would have received in the error-free case is called *lagging*. Flows that are neither leading nor lagging are called *in-sync*.

There are three important choices to make when defining an error compensation model. First, a procedure of lead/lag compensation must be selected. Second, how to determine the lead or lag of a traffic flow must be chosen. Third, a choice whether to prioritize keeping in-sync flows in the in-sync state must be made.

3.2.1 Lead/lag compensation

The compensation of leading and lagging flows can be classified into the following four categories:

No explicit compensation. The scheduler performs as in the error-free case, except that flows perceiving a channel error are skipped. No separate error-free reference model is used. Lead and lag values are not calculated, and the in-sync mode is not applicable.

The compensation, in the case of a WFQ scheduler, is the same as when the flow with minimum finish time receives access to the channel whenever it can transmit. A WRR scheduler simply skips the slots of a flow whose channel is in error. There is no explicit compensation for the lost slots. Fairness and throughput guarantees are met only for error-free flows.

The flow with maximum lag has channel access whenever it can transmit. There are two variants of this model:

- The flow with the *maximum lag* is allocated the channel whenever it can transmit. This compensation model requires the calculation of lead and lag (Section 3.2.2). When the flow with the largest lag among all flows becomes error-free, it can use the whole channel for its compensation.
- The flow with the *minimum finish time* has access to the channel whenever it can transmit. The compensation is done implicitly. A scheduling algorithm that selects the packet with the earliest finish time will give precedence to flows with lack of received service.

Flows being compensated for their lack of service will capture the whole channel and block the other flows. This is the fastest method of error compensation, and it provides long-term fairness for flows with bounded error. However, it does not comply with the other criteria of fair service in wireless networks.

Leading and lagging flows swap time slots. After a flow has been selected according the error-free model, compensation by slot swapping can be split into two phases: *lead compensation* (the decision whether or not to give up a slot), and *lag compensation* (the decision which flow may use the relinquished slot). The lag compensation is also used when a flow gives up its slot due to channel error.

The lead compensation has been narrowed to the following three possibilities:

- A leading flow *always* gives up its slots. The leading flow is blocked until it has lost all its lead. There is no short-term throughput guarantee for such a flow.
- A leading flow gives up a *constant fraction* of its slots. With this strategy, the short-term throughput bounds and fairness can be provided for error-free flows.
- A leading flow gives up a *varying fraction* of its slots, the fraction decreases exponentially as the lead is reduced [5]. This lead compensation provides a more graceful service degradation for leading flows. Short-term throughput and fairness guarantees can be met. However, it takes more time to compensate the lead of a flow.

The second phase of the slot swapping method consists of choosing a lagging flow to take over a time slot relinquished by another flow. The following three strategies are possible:

- The lagging flow with the *minimum finish time* is chosen. This lag compensation strategy requires a scheduling algorithm that uses the concept of finish times. Packets from lagging flows are scheduled in the same order as they would have been scheduled in the error-free case.
- The lagging flow with the *maximum lag* is chosen. Flows with a smaller amount of lag receive no compensation until the flow with the largest lag reaches an equal level of lag. Therefore, short-term fairness is not provided for these (now error-free) flows.
- The flow is chosen from a *Weighted Round Robin (WRR) allocation* of lagging flows. The flows within this WRR can be weighted either by their scheduling *weight*, or by their amount of *lag*. The compensation is distributed more evenly between lagging flows than in previous strategies. However, it may take longer to compensate a large lag if the WRR weights are proportional to the rate specification.

A fixed amount of bandwidth is reserved for compensation. A "compensation flow" with a fixed amount of bandwidth is declared in the scheduler [12]. This flow is used to compensate lost slots of flows with channel error, and it is scheduled in the error-free packet scheduler like any other flow. There is no independent error-free reference model to keep track of leading and lagging flows.

A flow that is selected to transmit a packet, but perceives a channel error, is skipped, and a reference to that flow (slot) is inserted into a data strucure at the compensation flow. When the compensation flow is selected for transmission, one of the flows whose reference is in the compensation queue can send a packet in its place. To be able to compensate all errors, it must be assured that enough bandwidth is reserved for compensation. The bandwidth reserved for compensation reduces the amount of bandwidth that can be reserved for other traffic flows.

The lag of each flow is monitored by the number of slots allocated for the flow in the compensation flow. Leading flows are not explicitly compensated since the compensation model does not keep track of the lead. Therefore, long-term fairness is not guaranteed.

3.2.2 Lead/lag mode

For compensation using slot-swapping between leading and lagging flows, and for the model in which the most lagging flow always receives channel access, the system needs to compute the lead and lag for each flow. Two ways of calculating these values are used in existing compensation models:

1. The lag of the flow is the difference between the error-free service and the real service received by the flow.
2. The lag of a flow is the number of slots allocated to the flow during which it could not send due to channel error, and another flow transmitted a packet in its place.

Lead is the negative value of lag. In lead/lag mode 2, a lack of service is only recognized as lag if another flow was able to use the time slot. Therefore, only the channel errors which caused another flow to receive more service will be compensated. Lead/lag mode 1 uses an independent error-free service model. This causes greater computing complexity. However, any differences between the error-free and the real service will be compensated.

3.2.3 In-sync mode

The in-sync mode was introduced in the WFS algorithm [5]. If a compensation algorithm uses this mode, in-sync flows will be kept in-sync if possible. In this mode, the order in which the algorithm searches for a flow to take over a relinquished slot is changed from *lagging–in-sync–leading* to *lagging–leading–in-sync*. If the in-sync mode is used in an algorithm, fewer flows will be leading, but these will accumulate a greater amount of lead.

4 A New Error Compensation Model for MPFQ

4.1 Error Compensation with Priority Classes

The MPFQ algorithm differs from most existing wireless scheduling algorithms since it is a prioritized algorithm with several classes scheduled under different policies. The individual scheduling policies (WFQ, WRR, and FIFO) are needed for different service requirements of each traffic class.

In MPFQ, the nature of the traffic in each priority class is known and can be taken into account for the design of the error compensation model. The shape of the incoming traffic and the delay requirements within each class are well defined. The new MPFQ compensation model takes advantage of these properties for each traffic class. It does not use a single compensation strategy, but rather compensates differently for each priority class. With this approach, an individual and therefore more desirable compensation for each traffic class can be achieved.

As the error-free traffic is scheduled in a prioritized manner in MPFQ, the error compensation will also be prioritized. Therefore, the traffic flows in a higher priority class (with tight delay guarantees) will be compensated with a higher priority than the traffic at lower priority classes (non-real-time). Error compensation for real-time flows should be performed as fast as possible. In non-real-time flows, fast compensation is not as important.

4.2 Lead/Lag Mode

In existing wireless scheduling algorithms, there are several ways of determining the lead or lag of a traffic flow. Since the rate requirements (and therefore the amount of service) are specified exactly for each flow, it is important to keep track of leading and lagging flows very accurately. This is done by using an independent error-free reference model. The lead or lag of the flows in each priority class is determined by comparing the amount of service in the error-free case with the amount of real service received.

4.3 CBR Compensation

The traffic category scheduled in the highest priority class in MPFQ is CBR traffic. The CBR traffic is scheduled using the WFQ algorithm with decoupling of bandwidth and delay. In this traffic class, the cell rate is well defined and has little variation. The delay bounds are tight. It is assumed that these properties are met by the incoming CBR traffic. However, the compensation model also needs to assure the same properties for the outgoing traffic. In the case, when channel errors are present, it is not always possible to guarantee small rate variations and tight cell delays. However, the compensation model should try to meet these requirements, while not disturbing other (error-free) flows in the CBR class.

Lead Compensation: A CBR flow with an error-free channel will never have many backlogged packets in the queue, since it is served at a well specified rate. Therefore, a CBR flow will not accumulate a great lead, and the lead compensation for CBR flows will not play a dominant role. However, for a practical implementation: A leading flow relinquishes a constant fraction of its allocated slots.

Lag Compensation: The lag compensation of CBR flows will use either bandwidth that is not reserved, time slots relinquished by other flows due to channel error, or time slots relinquished by leading flows in lower priority classes. A lagging CBR flow will be selected from a WRR allocation of lagging CBR flows. The weights in this WRR scheduler are proportional to the delay weights of the flows. Therefore, flows with tighter delay guarantees will be compensated faster than other CBR flows.

In-Sync Mode: A CBR flow will not usually become leading, as mentioned above. However, to increase the probability that an in-sync CBR flow will not be disturbed, the in-sync mode is enabled in this priority class.

4.4 Real-Time VBR Compensation

Rate changes are very common in VBR traffic. Since VBR traffic will be served by the MPFQ scheduler according to the Sustainable Cell Rate (SCR), it is assumed that there will sometimes be some packets backlogged, and that a VBR flow may become leading. The constraints on the output rt-VBR traffic are not as restricted as those of CBR traffic. The output rate of the rt-VBR traffic may vary, but the packet delay bounds remain tight.

Lead Compensation: Real-time VBR flows have tight delay bounds; lagging flows should therefore be compensated as fast as possible. However, a leading flow should not be completely blocked by the compensation of lagging flows, since its packets might be delayed longer than allowed. The lead compensation model chosen is a compromise between these two conditions. A leading flow will give up a constant fraction of its time slots.

Lag Compensation: Since the rate of a VBR flow is not constant, it is not appropriate to choose a lagging flow from a WRR allocation of lagging flows. For example, a lagging flow with a tight delay bound (a large weight in the WRR), but sending at a low rate at the moment, might gain an unfair advantage over another lagging flow sending at the Peak Cell Rate (PCR).

The lag compensation model chosen for the rt-VBR traffic class is to select the flow with the earliest finish time from among the lagging rt-VBR flows. This compensation model takes into account both delay weights, rate weights, and the rate at which the backlogged packets arrived.

In-Sync Mode: Packets of a rt-VBR flow that are sent at the PCR will usually be backlogged in the queue for some time even in the error-free case. The amount of service received by a rt-VBR flow will depend on the actual rate of the other rt-VBR flows. Therefore, it is not important for a rt-VBR flow to receive exactly the same amount of service (and not more service) as in the error-free case. The in-sync mode is not activated for this priority class.

4.5 Non-Real-Time VBR Compensation

The incoming non-real-time VBR traffic has similar characteristics to the real-time VBR traffic. Packets will sometimes be backlogged, even in the error-free case, and a nrt-VBR flow may become leading if other flows perceive channel errors. Packets in the nrt-VBR class have no tight delay bounds, so they may be blocked for some time if needed for the error compensation.

Lead Compensation: Since there are no tight delay bounds for nrt-VBR flows, a leading nrt-VBR flow can be blocked in order to let other flows compensate their lag. Since the rate of nrt-VBR flows may change, it is possible to reduce the rate to a minimum over short times. Blocking leading flows allows the fastest possible lag compensation for other flows.

Lag Compensation: In order to avoid having some flows that fall more and more behind their error-free service, while others can compensate their lag, the lagging flow with the largest lag is always chosen for compensation.

In-Sync Mode: The in-sync mode will not be used in this traffic class for the same reasons as in the rt-VBR class.

4.6 ABR and UBR Traffic

ABR traffic has a lag compensation as in the original MPFQ algorithm. It uses a common slot queue for all ABR flows, with the head-of-line slot recirculating in case of channel error. This may cause extra-delay in some cases, however, ABR traffic does not have tight delay specifications. Since the ABR traffic is in the second-lowest priority class, it is not necessary to do explicit lead compensation.

UBR traffic has no specified throughput or delay guarantees. It is scheduled in a best effort manner without any error compensation.

5 Simulations and Results

The MPFQ scheduler with the new MPFQ error compensation has been implemented in the *ns* simulator [6] along with an experimental class for simulating different error compensation strategies [2].

The theoretical improvements introduced with the new MPFQ compensation model were verified in simulations. In this section, the MPFQ algorithm with the new compensation model is compared to the MPFQ algorithm with non-explicit error compensation.

The important characteristics of fair error compensation (as explained in Section 3.1) are verified. These characteristics are long-term throughput bounds and fairness, delay bounds, and short-term throughput bounds and fairness.

5.1 Simulation Scenario

The scenario to compare the error compensation models consists of a base station with the MPFQ scheduler and three mobile hosts, each with a different distance to the base station. Each host is the destination of three flows whose parameters are given in Table 2.

Flow	CBR	rtVBR	nrtVBR
Peak Cell Rate	0.20	0.95	0.40
Sustainable Cell Rate	0.20	0.40	0.25
Maximum Burst Size	1	200	400
Priority Level	1	2	3

Table 2: Simulation traffic flows

Errors are simulated using a two-state Markov model for each mobile host. The Markov parameters are used as defined in [11] for the distances 10 ft, 90 ft, and 130 ft. One mobile host rarely has any errors, while the others suffer from errors during 2% and 8% of the simulation time respectively.

The duration of the simulations is 30 seconds (30000 packet time slots) for a channel bandwidth of 1 Mbit/s with a fixed packet length of 1000 bits. Different traffic and error patterns were used for a total of 20 simulations. The traffic and error data was read from trace files to run the simulations for the non-explicit compensation. The new MPFQ compensation model then was simulated under exactly the same conditions.

5.2 Long-Term Fairness

The long-term fairness was determined by counting the scheduled packets for each flow during the simulations. A difference between the number of scheduled packets in the error-free case and in the case with channel errors (using one of the compensation models) indicates unfairness in the long term.

No significant differences were observed. The largest deviation occurred at the nrt-VBR flow of the mobile host with the largest distance. This mobile received on the average 0.05% less service using the non-explicit compensation, and 0.04% less using the new MPFQ compensation model. Since these deviations are very small, the long-term fairness is provided by both compensations for this scenario.

Flow	Avg Delay		Avg WC Delay	
	non-expl.	new	non-expl.	new
CBR-10	0.407	0.394	9.958	8.702
CBR-90	0.823	0.822	65.155	65.330
CBR-130	2.023	2.045	90.830	92.056
rt-VBR-10	2.615	1.928	52.943	43.225
rt-VBR-90	3.538	3.066	85.892	81.438
rt-VBR-130	5.186	6.431	109.708	126.338

Table 3: Packet delays (milliseconds) for non-explicit compensation and for the new MPFQ compensation model

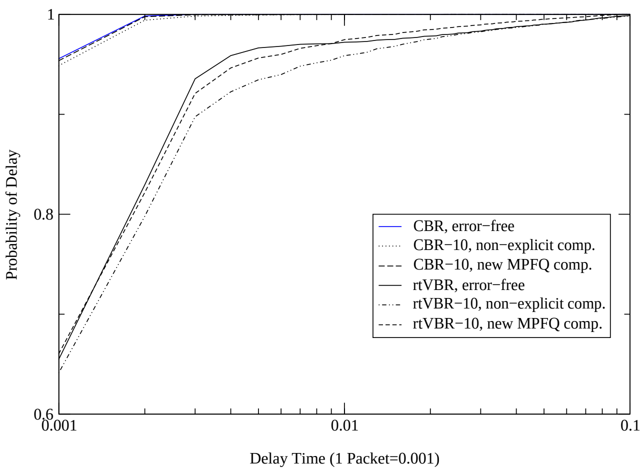

Figure 1: Cumulative Delay Distributions

5.3 Packet Delay

Channel errors have an important influence on packet delay. The average and worst-case delays become longer with higher error probabilities [8]. Since the non-explicit MPFQ error compensation uses the fastest possible way of compensation (use of the full bandwidth) in the two highest priority classes, no improvement for flows with channel errors can be expected.

However, a major drawback of the non-explicit compensation is that the flow separation is violated when error compensation occurs. The compensation of a flow blocks each of the other flows. Therefore, an improved delay distribution can be expected for flows with few channel errors.

Table 3 shows the average delay and average worst-case delay for the flows in the real-time priority classes. For the flows to the mobile host with 10 feet distance (low error rate), an improved delay is observed. In the case of the rt-VBR flow, the improvement is quite large. The new error compensation model improves the flow separation for real-time flows. However, flows having a higher error rate are delayed more by the new error compensation model since their compensation does not use the entire bandwidth.

Figure 1 shows cumulative packet delay distributions for the CBR and rt-VBR flows with 10 ft distance (almost error-free). The delay distributions from the simulation (both non-explicit compensation and the new MPFQ compensation) are compared with the distribution for the case where no flows are in error. The graph shows that the delay distribution for the new MPFQ compensation model is clearly closer to the error-free simulation than the delay distribution using non-explicit compensation. Error-free flows are disturbed less by the compensation of other flows in the new MPFQ compensation model.

5.4 Short-Term Fairness among Error-Free Flows

Short-term fairness among error-free flows can be shown by comparing the bandwidth of these flows at any time during the simulation. Within each of the three classes *leading*, *lagging*, and *in-sync*, the allocated bandwidth for each flow must be in proportion to the individual weights of the flows [10]. However, in order to compensate the lag of lagging flows, it is necessary that lagging flows use a higher amount of bandwidth than leading flows.

In a simulation scenario short-term fairness is difficult to show since all flows are not permanently backlogged. However, short-term fairness is maintained in the new MPFQ compensation model across all backlogged flows. This is accomplished by providing fair weighted service at any point in time.

6 Conclusion

In this paper, the design of a new MPFQ compensation model was presented that is based on well defined criteria for fair error compensation. This model uses a different compensation strategy for each priority class. Simulations showed that the new MFPQ error compensation model meets the long-term throughput guarantees and provides shorter delays for error-free flows in the presence of flows with channel errors (improved flow separation).

Next generation wireless networks will require QoS support in order to meet the varying traffic constraints of packet data. This support will rely heavily on a scheduler such as MPFQ that can provide QoS guarantees despite the unreliable underlying channel. In any such scheduler, the ability to properly compensate flows as they fade in and out of error is vital. The combination of this packet scheduler and compensation model will enable the widespread support of QoS measures on a wireless channel.

References

[1] ATM Forum. *Traffic Management Specification v4.0*, 1996. Document AF-TM-0056.000.

[2] S. Bucheli. MPFQ implementation for the *ns* simulator. `http://iwander.vlsi.uiuc.edu/wireless`.

[3] S. Bucheli. Compensation modeling for QoS in wireless networks. Diploma thesis, Swiss Federal Institute of Technology at Lausanne (EPFL) and University of Illinois at Urbana-Champaign, Feb. 2000.

[4] S. Lu, V. Bharghavan, and R. Srikant. Fair queueing in wireless packet networks. In *Proceedings of the ACM SIGCOMM Conference : Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM-97)*, pages 63–76, New York, Sept. 14–18 1997. ACM Press.

[5] S. Lu, T. Nandagopal, and V. Bharghavan. A wireless fair service algorithm for packet cellular networks. In *Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM-98)*, pages 10–20, New York, Oct. 25–30 1998. ACM Press.

[6] S. McCanne and S. Floyd. *ns* – Network Simulator. `http://www-mash.cs.berkeley.edu/ns`.

[7] J. R. Moorman and J. Lockwood. Implementation of the Multiclass Priority Fair Queuing algorithm for extending quality of service in existing backbones to wireless endpoints. In *GLOBECOM-99*, Dec. 1999.

[8] J. R. Moorman and J. Lockwood. Multiclass Priority Fair Queuing for hybrid wired/wireless quality of service support. In *ACM/IEEE WoWMoM/Mobicom 99*, Seattle, Aug. 20 1999. ACM Press.

[9] T. Nandagopal, S. Lu, and V. Bharghavan. A unified architecture for the design and evaluation of wireless fair queuing algorithms. In *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM-99)*, pages 132–142, Seattle, Aug. 15–20 1999. ACM Press.

[10] T. S. E. Ng, I. Stoica, and H. Zhang. Packet fair queueing algorithms for wireless networks with location-dependent errors. In *IEEE INFOCOM-98*, Mar. 1998.

[11] G. T. Nguyen, R. H. Katz, B. Noble, and M. Satyanarayanan. A trace-based approach for modeling wireless channel behavior. In *Proceedings of the Winter Simulation Conference*, pages 597–604, Dec. 1996.

[12] P. Ramanathan and P. Agrawal. Adapting packet fair queueing algorithms to wireless networks. In *Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM-98)*, pages 1–9, New York, Oct. 25–30 1998. ACM Press.