A QoS Framework for Stabilized Collision Channels with Multiuser Detection

Kristoffer Bruvold, Raghuraman Mudumbai and Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106, USA
Email: {bruvold,raghu,madhow}@ece.ucsb.edu

Abstract— **Recent work has shown that cross-layer optimization of the physical layer and Medium Access Control for a wireless collision channel, based on a receiver with adaptive multiuser detection capability, is capable of providing significantly better performance than classical Aloha. The basic features of such a system are multipacket reception (MPR) capability, and the ability (with high probability) to estimate the number of contending users even when the packets are not successfully received. We provide an analytical model that includes these features, and use it to derive methods for backlog estimation and stabilization. Two classes of users are considered: high priority users with Quality of Service (QoS) requirements, who must succeed within a deadline with a specified probability; and low priority users whose throughput we wish to maximize, while maintaining the QoS for high priority users, and keeping the overall system stable. We obtain contention policies that ensure QoS and stability, based on backlog estimates obtained by extending Rivest's pseudo-Bayesian technique for classical Aloha. The channel throughput and the achievable QoS is characterized as a function of the arrival rates for high and low priority users. Finally, we apply these methods to simulations of a system employing Differential Minimum Mean Squared Error (DMMSE) adaptive multiuser detection, and find that the analytical model provides accurate guidelines for design and performance predictions.**

I. Introduction

The classical slotted Aloha model [1]–[3] is based on the assumption that exactly one packet can be successful in a given slot. Advances in multiuser detection techniques over the last two decades imply, however, that systems with multipacket reception capabilities are now becoming practically feasible. In recent work [4], we have shown that joint optimization of the physical and MAC layers based on a receiver capable of adaptive multiuser detection leads to an Aloha-like system which is capable of supporting both high priority users with deadline constraints, and low priority users with best effort service. The specific application considered in [4] was the design of a reservation channel for rapid handoffs in a *pseudocellular* network, which supports vehicular mobility, as in cellular networks, using small pseudocells covered by a wireless local area network (WLAN) type infrastructure. For example, a voice connection at vehicular speeds can be supported by ensuring that the mobile can make a timely reservation with the AP whose pseudocell it is entering, thus achieving a mobile-centric handoff. QoS for pseudocellular networks was also considered in [5], where contention policies are derived to provide support for deadlines without MUD capabilities at the receiver. MUD capabilities at the receiver enable us to achieve much higher throughput in the pseudocellular reservation channel, as was demonstrated in [4]. However, the study of collision channels with MPR is of fundamental importance beyond its application to pseudocellular networks. The purpose of this paper is to abstract an analytical model of the system in [4], in order to obtain, under realistic constraints, a general design framework for stabilization and QoS for a collision channel with MPR capability. In contrast to prior work on Aloha with MPR, our modeling assumptions are specifically guided by the capabilities of the adaptive multiuser detection strategy used at the receiver.

Uncontrolled Aloha with MPR was studied in [6], where it was shown that that the system is stable for arrival rates below a threshold if the number of successes has a nonzero limit as the collision size gets large. The latter condition amounts to requiring some form of capture, and is typically not satisfied when the number of contending users exceeds the capability of the multiuser detector (in which case, most likely, all users fail to decode). There are also other similarly restrictive stability conditions that have been considered in the literature, such as assuming a lower bound on the signal-to-interference ratio [7]. In a wireless network with mobiles contending for communicating with an access point (AP), however, the AP can control the contention policies used by the mobiles based on its knowledge of the network conditions. It is known that classical Aloha can be stabilized by controlling the transmission probability based on estimates of the backlog [8], [9]. In this paper, we extend the pseudo-Bayesian backlog estimation method in [8] to take advantage of the receiver's capabilities for MPR and for providing feedback on the collision size, and apply this to multiple priority classes, including both delay-constrained and delay-tolerant traffic. We then employ these estimates to design stable contention policies that provide QoS to *Hi*-priority users (i.e., assuring that they meet a deadline with a specified probability), while maximizing the throughput for *Lo*-priority users. We develop a multi-dimensional Markov chain description that specifies the throughput as a function of the traffic arrival rates for each

This work was supported by the National Science Foundation under grants EIA-0080134 and ANI-0220118 (ITR), and by the Office Naval Research under grant N00014-03-1-0090.

priority class. With the throughput characteristics obtained from the Markov chain analysis, the backlog estimates are used to generate the contention policies (e.g., to be broadcast by the AP in its beacon prior to each frame). Note that the availability of estimates of collision size from the physical layer simplifies the design, relative to prior work which relies on higher layer information for backlog estimation [10].

We assume only three general properties of the receiver for the analytical model of the MPR channel. First, although users transmit simultaneously in synchronized contention frames, there is a limit, N, to the maximum number of users that can be successfully decoded by the receiver. If the number of transmissions is above this limit, none can be decoded. As an example, this property is analogous to the DS-CDMA processing gain limit on the MUD capabilities of the system in [4]. Second, there is a limit, P, to the number of unique ways, termed *virtual subslots*, that users may use to transmit their packets to the AP. When transmitting a packet, a user randomly chooses one of these virtual subslots and, so long as the total number of transmissions is less than N and no other user has chosen the same virtual subslot, that user's packet is successfully decoded. The virtual subslots may be provided by any novel receiver technique, an example of which is the user's choice of training symbol sequence in [4]. The third assumption, which is critical to our backlog estimation algorithms, is that the receiver is able to determine the total number of transmissions that occurred in a given frame, even when some or none of them are successfully decoded. This assumption is approximately satisfied by the DMMSE receiver employed in our previous work [4], [11].

The remainder of this paper is organized as follows. In Section II, we present the throughput analysis for the generalized model with $P > N$ which is used to design the contention policies in Section III. In Section IV, we describe the extension of Rivest's backlog estimation technique [8] to the MPR case. Section V compares the analytical results with the simulated performance of a DMMSE receiver implementation [4]. Section VI provides concluding remarks.

II. Throughput Analysis

We consider here a simple special case of the MPR multi-access contention channel in [6]: a user may collide and fail only if either the total number of users exceeds a threshold, N, or another user chooses the same virtual subslot from the P available. For this channel, we extend classical ALOHA theory [12] to write the expression for the throughput as:

$$T(G) = \sum_{k=1}^{N} \overline{\mathrm{N}}_k(P) P_G(k). \tag{1}$$

Here, $P_G(k) = \frac{1}{k!}G^k e^{-G}$ is the rate G Poisson probability of k arrivals. The expected number of successes given k users choosing from P virtual subslots is given by $\overline{\mathrm{N}}_k(P) = \sum_{n=1}^{k} nF(n|k)$, where $F(n|k)$ is the probability of n out of k users choosing unique virtual subslots. Using combinatorics, it can be shown that

$$\sum_{j=n}^{k} \binom{j}{n} F(j|k) = \binom{k}{n} \frac{P!(P-n)^{k-n}}{P^k (P-n)!}, \tag{2}$$

with the resulting set of linear equations recursively solvable for $F(n|k)$.

Since, for $N < P$, $T(G)$ and $G_{opt} \triangleq \arg\max_G T(G)$ are increasing functions of N, and $G_{opt} \le N$, rates higher than N will produce decreasing throughput. Therefore, we restrict the analysis to the range $G \in [0, N]$.

We consider now the problem of providing QoS support for a single class of *Hi*-priority users where each needs to make a successful transmission within a delay deadline of D reservation frames with probability of at least R_s. The expectation is that the supported Poisson arrival rates will be less than G_{opt} in order to satisfy this QoS. For the simple case of $D = 1$, each user either succeeds on its first attempt, or its deadline expires. Thus, there are no retransmissions and the total transmission process is the same as the arrival process of new users, which is Poisson with rate $G = \lambda$. Further, the probability of the deadline expiring is $P_{exp} = 1 - \frac{T(G)}{G}$ and is completely determined by (1). As an example, for $N = 11$ and $P = 31$, $\lambda = 0.311$ is the maximum sustainable arrival rate for which $P_{exp} \le 1 - R_s = 0.01$. This is much smaller than the corresponding $G_{opt} = 7.757$ with throughput $T(7.757) = 5.20$. Intuitively, if the QoS requirement is relaxed by increasing the delay deadline to $D = 3$, a higher arrival rate might be supported. However, the transmission process then consists of a combination of a Poisson arrival process and a retransmission process, and the total process is no longer Poisson. Therefore, a Markov model is used for analysis. Before deriving the model, we state the assumptions and some notation.

- Users contend with a single packet in the MPR channel, and new arrivals form a Poisson process with rate λ. Without loss of generality, we assume that $N = 11$, $P = 31$, $D = 3$, and $R_s = 0.99$.
- Each *Hi*-priority user transmits in each frame with probability $p_{Tx} = p_{Hi} = 1$.
- For $D = 3$, the contention channel backlog is modeled as a two-dimensional Markov process where the i^{th} state $S_i = (\mathrm{N}_1, \mathrm{N}_2)$ represents the number of backlogged users with N_j having already collided j times. Since $p_{Tx} = 1$, all packets are lost for any state where $\sum_j \mathrm{N}_j > N$. Therefore, the state space is truncated by restricting N_j to the range $[0, N+1]$. This effectively neglects the probability of arrival bursts larger than $N + 1$, however, for the range of λ satisfying $R_s \ge 0.99$, this assumption is justified. The total number of states in the Markov chain is $L = (N+2)^{D-1}$.
- For a system in state $S_i = (\mathrm{N}_1, \mathrm{N}_2)$ at the beginning of a frame where there are k new arrivals, let n_1, n_2, and n_k denote the number of successes from the N_1, N_2, and k users, respectively. Further, the general probability of getting the a-tuple of $(n_1, n_2, \ldots, n_a)$ successes from

the corresponding a-tuple of $(\mathrm{N}_1, \mathrm{N}_2, \ldots, \mathrm{N}_a)$ users in a frame is denoted by $f(n_1, n_2, \ldots, n_a | \mathrm{N}_1, \mathrm{N}_2, \ldots, \mathrm{N}_a)$.

In order to find the probability of deadline expiration, P_{exp}, we compute the stationary probability distribution $\mathbf{p} = [p_1, p_2, \ldots, p_L]^T$, where $p_i = p(S_i)$, the stationary probability of state S_i. The probability of transition from state $S_i = (\mathrm{N}_1, \mathrm{N}_2)$ to state $S_j = (k - n_k, \mathrm{N}_1 - n_1)$ is given by

$$Q_{i,j} = \sum_{k=0}^{N+1} P_\lambda(k) Q_{i,j|k}, \tag{3}$$

where $P_\lambda(k)$ is the (Poisson) probability of k new arrivals. The conditional probability of transition from state S_i to state S_j given k new arrivals, $Q_{i,j|k}$, is given by

$$\begin{aligned} Q_{i,j|k} &= \sum_{n_2=0}^{\mathrm{N}_2} f(n_1, n_2, n_k | \mathrm{N}_1, \mathrm{N}_2, k) \qquad (4) \\ &= \sum_{n_2=0}^{\mathrm{N}_2} \frac{\binom{\mathrm{N}_1}{n_1}\binom{\mathrm{N}_2}{n_2}\binom{k}{n_k}}{\binom{\mathrm{N}_1+\mathrm{N}_2+k}{n_1+n_2+n_k}} F(n_1 + n_2 + n_k | \mathrm{N}_1 + \mathrm{N}_2 + k). \end{aligned}$$

Using (3), (4), and the results from (2), we can compute the stationary probability distribution $\mathbf{p}$, and similarly $E[\mathrm{N}_{exp,i} = \mathrm{N}_2 - n_2]$, the expected number of expired users in state S_i. Averaging $E[\mathrm{N}_{exp,i}]$ over $\mathbf{p}$ yields the overall expected number of expirations, $E[\mathrm{N}_{exp}]$. Finally, $P_{exp} = \frac{1}{\lambda} E[\mathrm{N}_{exp}]$. The state space in Fig. 1 illustrates the contribution of each state to the overall expiration rate.

We consider now the case of supporting best effort traffic, using the "spare capacity" from when the *Hi*-priority users' arrival rate, λ_{Hi}, is less than the maximum sustainable for QoS, computed as previously outlined. The idea is that depending on λ_{Hi}, it should be possible to accommodate a certain amount of best effort traffic with rate λ_{Lo} without sacrificing *Hi* QoS. In this situation, the resulting channel utilization should be higher than when all users have QoS constraints, but lower than $T(G_{opt})$ for no QoS constraints on any users.

We extend the previous model by including the extra *Lo*-priority arrival process, which changes the transition probabilities of the Markov model. Omitting the details for lack of space, (3) and (4) can be rewritten for this case as:

$$Q_{i,j} = \sum_{k_{Hi}, k_{Lo}=0}^{N+1} P_{\lambda_{Hi}}(k_{Hi}) P_{\lambda_{Lo}}(k_{Lo}) Q_{i,j|k_{Hi},k_{Lo}}, \tag{5}$$

$$Q_{i,j|k_{Hi},k_{Lo}} = \sum_{n_2=0}^{\mathrm{N}_2} f(n_1, n_2, n_{Hi}, n_{Lo} | \mathrm{N}_1, \mathrm{N}_2, k_{Hi}, k_{Lo}). \tag{6}$$

As before, we compute $\mathbf{p}$ and P_{exp}. Results for the cases $\lambda_{Hi} = \{2.0, 2.5, 3.0\}$ are shown in Fig. 2. From these results, the maximum sustainable λ_{Lo} corresponding to a given value of λ_{Hi}, i.e., the *Lo* capacity, $C_{Lo}(\lambda_{Hi})$, is determined. Next we consider the design of contention policies that maintain these throughputs and QoS for the multiple priority class traffic.

III. Contention Policies

We extend the work of Rivest in [8] on the single-packet, single priority class channel to the multi-packet, multiple priority class channel in our model. For Rivest's model [8], the channel capacity was $C = 1$ and the optimal transmission policy, which minimizes the average backlog while maintaining the maximum throughput, was shown to be

$$p_{opt}[n] = \min\left(1, \frac{C}{\mathrm{N}[n]}\right) = \min\left(1, \frac{1}{\mathrm{N}[n]}\right), \tag{7}$$

for the backlog $\mathrm{N}[n]$ at the beginning of frame n. The resulting expected number of contention attempts in the frame is

$$E\left[\mathrm{N}_{att}[n]\right] = p_{opt}[n]\mathrm{N}[n] = C = 1. \tag{8}$$

For large $\mathrm{N}[n]$, this process is well approximated as a Poisson arrival process with rate $C = 1$ [13].

To extend (7) to the multipacket, multiple priority class system, we use $C_{Lo}(\lambda_{Hi})$, the maximum *Lo*-priority capacity as a function of λ_{Hi}, as prescribed in Section II. This is used to scale (7), resulting in the optimal *Lo* transmission policy given by

$$p_{Lo}[n] = \frac{C_{Lo}(\lambda_{Hi})}{\mathrm{N}_{Lo}[n]}, \quad 0 \le p_{Lo}[n] \le 1, \tag{9}$$

where $\mathrm{N}_{Lo}[n]$ is the *Lo* backlog at the beginning of frame n. The expected number of contending *Lo* users scales to

$$E\left[\mathrm{N}_{att,Lo}[n]\right] = C_{Lo}(\lambda_{Hi}), \tag{10}$$

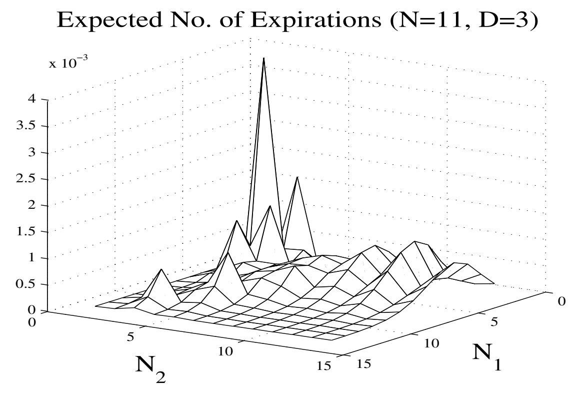

Fig. 1. Expected number of deadline expirations in each state.

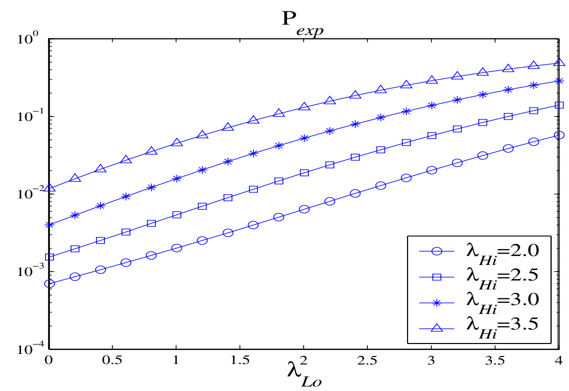

Fig. 2. QoS for *Hi* priority users as a function of *Lo* priority traffic in an uncontrolled system.

forming a Poisson process of rate $C_{Lo}(\lambda_{Hi})$ for large $\mathrm{N}_{Lo}[n]$. Thus, if a reasonably accurate estimate of $\mathrm{N}_{Lo}[n]$ is available, the throughput for the *Lo* users can be maximized without violating the QoS guarantee of the *Hi* users.

In our experiments, a slightly different form of (9) has shown to offer the same or better performance. In the alternate form, the "instantaneous" *Lo* capacity, $C_{Lo}(\mathrm{N}_{Hi}[n])$ is used in the numerator of (9) instead of the capacity, $C_{Lo}(\lambda_{Hi})$, based on the average arrival rate λ_{Lo}. There are two potential advantages to this modification. First, if $\mathrm{N}_{Hi}[n] < \lambda_{Hi}$, there is additional capacity available for use by *Lo* users, enhancing their throughput. Second, when $\mathrm{N}_{Hi} > \lambda_{Hi}$, $E\left[\mathrm{N}_{att,Lo}[n]\right]$ can be scaled back to allow additional capacity to serve the burst of *Hi* users at the highest possible QoS.

IV. Backlog Estimation

For the contention policies prescribed in Section III, a technique is required for estimating the backlog in an uncoordinated fashion. To do this, we again extend the framework in [8] to apply to our higher capacity system, and then incorporate the additional information available at the receiver regarding the previous contention frame's outcome. In [8], the estimated number of backlogged users at the start of a frame is updated based on three possible outcomes in the previous frame: a *hole*, when no users transmit; a *success*, when only one user transmits; and a *collision*, when more than one user transmits. Given the expected number of new arrivals, λ, and the outcome in frame n, the pseudo-Bayesian estimate for the number of backlogged users in frame $n+1$ is computed via

$$\hat{\mathrm{N}}[n+1] = \hat{\mathrm{N}}[n] + c(f[n]) + \lambda, \tag{11}$$

where $f[n] = \{0, 1, e\}$ is the feedback for a hole, success, or collision, respectively, in frame n and

$$c(f[n]) = \begin{cases} -1 & \text{for } f[n] = \{0,1\}, \text{ and} \\ \frac{1}{e-2} & \text{for } f[n] = e. \end{cases} \tag{12}$$

Thus, the backlog estimate for the next frame is the previous estimate adjusted by a correction factor for the outcome in the previous frame plus the expected number of new arrivals.

A. Multipacket, Multiple Priority Backlog Estimation

In our model, we assume that the receiver is capable of determining how many users transmitted packets in a given frame, regardless of the number who succeeded or failed (in fact, we have demonstrated a receiver capable of this detection [11] so long as the total number contending is below a threshold of roughly $2N$, where N is the DS-CDMA processing gain). Further, since all three types of outcomes–holes, successes and collisions–are possible in the same slot, (12) must be modified and it must also account for multiple priority classes. Therefore, a backlog estimate must be maintained separately for each priority class (thus, any cost associated with fine-tuning the algorithm as in Section III is already mandated).

Since the numbers of holes or collisions in a reservation slot are not necessarily equal to 1, we must define these quantities based on the information at the receiver. The number of collisions is given by

$$\mathrm{N}_{coll}[n] = \mathrm{N}_{att}[n] - \mathrm{N}_{succ}[n], \tag{13}$$

where $\mathrm{N}_{succ}[n] = \mathrm{N}_{succ,Hi}[n] + \mathrm{N}_{succ,Lo}[n]$. The definition in (13) also allows for the case when users might individually fail because their packet bit error rate (BER) is too high due to interference. The number of holes is the difference between the expected number of attempts and the actual number of attempts in a frame. As in [8], this intuitively represents the unused portion of the capacity given the contention policy in (9) is based on backlog estimates:

$$\mathrm{N}_{hole}[n] = \max\Big(\big(\hat{\mathrm{N}}_{Hi}[n] + C_{Lo}(\hat{\mathrm{N}}_{Hi}[n])\big) - \mathrm{N}_{att}[n], 0\Big). \tag{14}$$

With these definitions, the general form of the correction factor in (12) for priority class i becomes:

$$c_i[n+1] = -\mathrm{N}_{succ,i}[n] - \hat{\mathrm{N}}_{hole,i}[n] + \hat{\mathrm{N}}_{coll,i}[n] + \hat{\mathrm{N}}_{corr,i}[n], \tag{15}$$

where $\hat{\mathrm{N}}_{hole,i}[n]$ and $\hat{\mathrm{N}}_{coll,i}[n]$, are the estimated hole and collision measures, respectively, and $\mathrm{N}_{succ,i}[n]$ is known. $\hat{\mathrm{N}}_{corr,i}[n]$ is an additional correction factor explained later in detail.

Although we assume the receiver detects $\mathrm{N}_{coll}[n]$, we do not assume that it can distinguish the priority classes of the users who collided. Therefore, when accounting for failures and holes in the backlog estimates, a proportion factor, $\hat{\beta}_i[n] = p_i[n]\hat{\mathrm{N}}_i[n] / \sum_k p_k[n]\hat{\mathrm{N}}_k[n]$, representing the expected proportion of class i users from all users contending in frame n, is computed based on the backlog estimates, $\hat{\mathrm{N}}_k[n]$, for each class. For the *Hi* and *Lo* priority classes used here, these factors are given by

$$\hat{\beta}_{Hi}[n] = \frac{p_{Hi}[n]\hat{\mathrm{N}}_{Hi}[n]}{p_{Hi}[n]\hat{\mathrm{N}}_{Hi}[n] + p_{Lo}[n]\hat{\mathrm{N}}_{Lo}[n]}, \tag{16}$$

$$\hat{\beta}_{Lo}[n] = 1 - \hat{\beta}_{Hi}[n]. \tag{17}$$

The proportion factors are then applied to generate the estimated hole and failure measures for class i:

$$\hat{\mathrm{N}}_{coll,i}[n] = \hat{\beta}_i[n]\mathrm{N}_{coll}[n], \quad \hat{\mathrm{N}}_{hole,i}[n] = \hat{\beta}_i[n]\mathrm{N}_{hole}[n]. \tag{18}$$

As described next, in addition to $\hat{\mathrm{N}}_{corr,i}[n]$, these measures are applied conditionally to (15).

In (15), the primary function of $\hat{\mathrm{N}}_{corr,i}[n]$ is to account for the *Hi* users' effect on both $\hat{\mathrm{N}}_{Hi}[n]$ and $\hat{\mathrm{N}}_{Lo}[n]$. If a *Hi* user succeeds after having collided at least once, these extra collisions are removed from the backlog estimates $\hat{\mathrm{N}}_{Hi}[n]$ and $\hat{\mathrm{N}}_{Lo}[n]$ retroactively, where they were included in (18). Also, *Hi* users whose delay budgets have been exceeded and are exiting the system are removed from the *Hi* backlog estimate. Finally, since *Hi* users contend with $p_{Hi} = 1$, if $\mathrm{N}_{coll}[n] = 0$, $\hat{\mathrm{N}}_{Hi}[n+1] = \lambda_{Hi}$. Therefore, any *Hi* failure measures for the previous $D-1$ frames are retroactively attributed instead to the *Lo* backlog estimate.

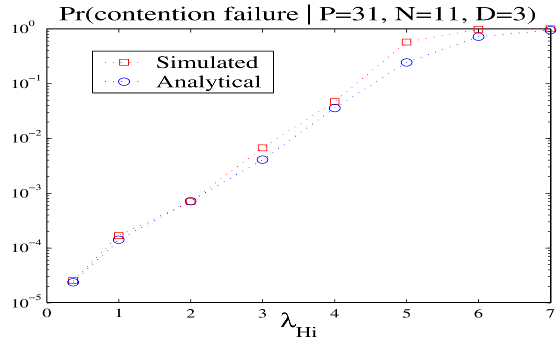

Fig. 3. Probability of contention failure for a delay budget of $D = 3$ frames as a function of λ_{Hi}, with only *Hi* priority traffic contending.

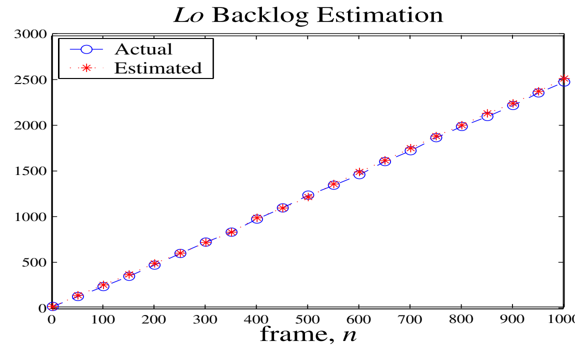

Fig. 4. Estimated and actual *Lo* backlog with both *Hi* and *Lo* priority traffic contending.

V. Numerical Results

We now compare the results of our analysis with simulation results for our jointly optimized PHY/MAC design in [4] based on the DMMSE receiver [11]. The DS-CDMA system in [4] has processing gain $N = 11$, and 124 training symbols, values chosen to be comparable to 1-2 Mbps 802.11b WLANs [14] (however, the 802.11b format is not amenable to multipacket reception). The number of virtual subslots, $P = 31$, equals the number of distinct training sequences that transmitting users can choose from. Each user experiences flat Rayleigh fading. The QoS required for *Hi* users is $R_s \geq 0.99$.

First, the analytical model for the throughput for the case where only *Hi* users contend with $D = 3$ is compared with the actual system. Fig. 3 shows the analytical and simulated curves for the expiration rates in the QoS region of interest. The close match observed therein validates the analytical model. Further, the system is shown to be capable of supporting up to $\lambda_{Hi} = 3.5$ arrivals per frame at an expiration rate of $P_{exp} = 1 - R_s \leq 0.01$. As expected, this rate is between the achievable $\lambda_{Hi} = 0.311$ for the stringent constraint of $D = 1$ and the maximum achievable $T(G_{opt}) = 5.20$ from (1) without QoS.

Finally, Fig. 4 shows a trace over 1000 frames of the backlog estimates for *Lo* priority users compared with the actual backlog values. For these curves, the value of the estimate was used when generating the contention policy from (9). The close match of the curves show that the technique provides accurate estimates that are well suited for dynamic control of the contention policy.

VI. Conclusion

While earlier work had shown that the performance of Aloha can be improved with MPR capability provided by multiuser detection, a novel feature of our work is the exploitation of collision size estimates provided by the receiver. The latter is what enables us to obtain accurate backlog estimates for the different priority classes, which in turn are used to control the system to operate in the desired QoS/stability regime. Our results show the large performance gains to be obtained from cross-layer optimization. Moreover, the close match between our analytical model and the simulation model with the DMMSE receiver shows that our analysis is accurate and tractable, and offers a useful abstraction of the cross-layer interactions in the architecture.

References

[1] N. Abramson, "The ALOHA system - another alternative for computer communications," in *Proc. Fall 1970 AFIPS Computer Conference*, 1970, pp. 281–285.

[2] L. G. Roberts, "ALOHA packet system with and without slots and capture," in *ASS Note 8*, Stanford Res. Inst., Adv. Res. Projects Agency, Network Inform. Ctr., Stanford, CA, Jan. 1972.

[3] L. Kleinrock and S. S. Lam, "Packet switching in a multiaccess broadcast channel: Performance evaluation," *IEEE Trans. Commun.*, vol. 23, no. 4, pp. 410–423, Apr. 1975.

[4] K. Bruvold and U. Madhow, "Adaptive multiuser detection for mobile-centric fast handoffs in pseudocellular wireless networks," in *Proc. 58th IEEE Vehicular Technology Conference (VTC Fall'03)*, Orlando, FL, Oct. 6–9, 2003.

[5] R. Mudumbai, G. Barriac, and U. Madhow, "Optimizing medium access control for rapid handoffs in pseudocellular networks," in *Proc. 60th IEEE Vehicular Technology Conference (VTC Fall'04)*, Los Angeles, CA, Sept. 26–29, 2004.

[6] S. Ghez, S. Verdú, and S. C. Schwartz, "Stability properties of slotted aloha with multipacket reception capability," *IEEE Trans. Automat. Contr.*, vol. 33, no. 7, pp. 640–648, July 1988.

[7] V. Naware and L. Tong, "Stability of slotted ALOHA with spatial diversity," in *Proc. 2003 IEEE International Conference on Communications (ICC'03)*, Anchorage, AK, May 11–15, 2003, pp. 463–467.

[8] R. L. Rivest, "Network control by bayesian broadcast," *IEEE Trans. Inform. Theory*, vol. 33, no. 3, pp. 323–328, May 1987.

[9] B. Hajek and T. van Loon, "Decentralized dynamic control of a multiaccess broadcast channel," *IEEE Trans. Automat. Contr.*, vol. 27, pp. 559–569, 1982.

[10] Q. Zhao and L. Tong, "A multiqueue service room MAC protocol for wireless networks with multipacket reception," *IEEE/ACM Trans. Networking*, vol. 11, no. 1, pp. 125–137, Feb. 2003.

[11] U. Madhow, K. Bruvold, and L. J. Zhu, "Differential MMSE: A framework for robust adaptive interference suppression for direct sequence CDMA over fading channels," *submitted for publication*.

[12] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 1992.

[13] H. Stark and J. W. Woods, *Probability, Random Processes, And Estimation Theory For Engineers*, 2nd ed. Upper Saddle River, NJ: Prentice-Hall, 1994.

[14] "Supplement to IEEE standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements. part 11: wireless LAN medium access control (MAC) and physical layer (PHY)," IEEE, Tech. Rep. IEEE Std. 802.11b-1999, 2000.