Achievable throughput and queueing delay for imperfect cooperative retransmission MAC protocols

Steven Weber and Athina Petropulu
Drexel University, Dept. of ECE, Philadelphia, PA 19104

***Abstract*— Cooperative retransmission (CR) for cellular uplink employs recovery epochs where colliding packets are retransmitted until all packets are recoverable by the base station. We consider a general class of queueing dynamics, appropriate for cellular uplink MAC protocols, where time slots are grouped into epochs, and each epoch results in some subset of the active nodes completing a transmission to the base station. Employing moment generating functions, we express the queue length in terms of the number of arrivals during each epoch, and from here we obtain general expressions for throughput and delay (queueing plus service). We specialize these general results for four MAC uplink protocols: $i)$ independent queues with no collisions (providing an upper bound on performance), $ii)$ slotted Aloha, $iii)$ perfect cooperative retransmission (each time slot yields a new independent signal combination), and $iv)$ imperfect cooperative retransmission (a time slot gives useful information with probability $1-q$).**

I. Introduction

Cooperative retransmission for cellular uplink makes profitable use of colliding transmission attempts that would otherwise be wasted. Upon a collision of $j \geq 2$ uplink transmission attempts by j nodes in a time slot, a base station employing a cooperative retransmission MAC will initiate a recovery period. The recovery period consists of a set of j cooperating nodes retransmitting either their own packets, if they were involved in the collision, or the signal that they heard during the collision slot. All nodes not involved in the recovery refrain from transmission until the completion of the recovery period, announced by the base station, after which all nodes in the cell are free to contend for channel access. The base station stores these repeated collided signals, and will be able to recover all j packets as soon as j linearly independent collided signals are available. The independence of the collided signals is achieved by exploiting some form of diversity in the environment. Various protocols have been proposed for leveraging this diversity. The NDMA protocol of [1] exploits temporal variation in channel fading gains, while the Alliances protocol of [2] exploits spatial diversity. The recovery period requires $k \geq j$ time slots, where k is called the *duration* of the recovery.

Perfect recovery occurs when the duration of the recovery (k) equals the order (j), *i.e.*, each time slot produces a new independent collided signal at the base station. In such a case the cooperative retransmission protocol achieves perfect throughput for the recovery: j packets received in j time slots. *Imperfect* recovery occurs when the recover duration exceeds the order: $k > j$; in such a case the protocol achieves an imperfect throughput of j packets in k time slots. We assume there is a fixed probability q that each time slot will not generate an independent mixture, and thus *on average* a collision of order j requires $j/(1-q)$ time slots to resolve.

In §II we consider a general model of queueing dynamics appropriate for cellular uplink MAC protocols. In this general model time slots are grouped into epochs, where each epoch is either successful or unsuccessful for each node j, depending upon whether or not the base station receives a packet from node j at the end of the epoch. We find the moment generating function (mgf) for the steady state queue length in terms of the mgfs for the number of arrivals during both successful and unsuccessful epochs. From this characterization we obtain expressions for throughput and delay. In §III we specialize the general results from §II to four specific cellular uplink protocols: $i)$ independent queues with no collision (providing an upper bound on performance), $ii)$ slotted Aloha, $iii)$ perfect cooperative retransmission (where each time slot yields a new independent signal combination), and $iv)$ imperfect cooperative retransmission (where a time slot gives useful information with probability $1-q$). In each case we characterize the expected throughput and overall (queueing plus service) delay as a function of the contention probability p_c. The contention probability p_c is the probability that each node (with a packet to send) elects to attempt a transmission in each time slot that is not part of a recovery epoch. In §IV we provide some numerical results that illustrate the performance of the four protocols, and we offer a brief conclusion in §V.

The technical approach behind the results in §2 in this paper is inspired by §V-C in [1], but there are several innovations above and beyond that work. In particular, [1] restricts attention to perfect recovery ($q = 0$) without random contention ($p_c = 1$) for Poisson arrivals ($\sigma^2 = \lambda$). Even more, they do not derive the delay equation directly from the queueing dynamics as we do. Compare (5) in this paper with (41) in [1], where their delay expression is obtained from an $M/G/1$ queue with vacations. This paper is also similar in spirit to our own related work in [3] on slotted Aloha and [4] on imperfect cooperative retransmission protocols. The key distinction between [3] and this work is that the former emphasizes the independence assumption required to gain traction on the evolution of the queueing dynamics, while the key distinction between [4] and this work is that the former studies performance under a saturation assumption when queues are always backlogged. This means the natural delay metric in [4] is *service* delay, while in this paper we study queueing plus service delay.

This work is supported by the Office of Naval Research (Grant N00014-07-1-0500). Corresponding author: Steven Weber, sweber@ece.drexel.edu.

II. A GENERAL CLASS OF QUEUEING DYNAMICS

We consider the following generic model for a cellular uplink protocol:

- The cell consists of a set of $\mathcal{J}$ nodes ($|\mathcal{J}| = J$). Each node has a stochastic arrival process of new packets to be transmitted to the base station. The arrival process is independent across nodes, and independent of the state of the system. The arrival process is characterized by the mean (λ) and variance (σ^2) of the number of arrivals per time slot.
- Time is slotted and all transmissions are synchronous, lasting for one time slot. Time slots are grouped into epochs, and epochs (not slots) are indexed by t.
- Epochs are individually labeled as successful or unsuccessful for each node j: epoch t is successful if at the end of the epoch the base station has successfully decoded a packet from node j. Else, epoch t is unsuccessful for node j.

The queueing dynamics for each node j under this model are:

$$Q^j(t+1) = Q^j(t) + R_s^j(t)\mathbf{1}_{\mathcal{S}^j(t)} + R_u^j(t)\mathbf{1}_{\mathcal{U}^j(t)} - \mathbf{1}_{\mathcal{S}^j(t)}. \quad (1)$$

Here:

- $Q^j(t)$ is the queue length at node j at the start of epoch t;
- $R_s^j(t)$ is the number of arrivals at node j during *successful* epoch t;
- $R_u^j(t)$ is the number of arrivals at node j during *unsuccessful* epoch t;
- $\mathcal{S}^j(t)$ is the event that epoch t is *successful* for node j;
- $\mathcal{U}^j(t) = (\mathcal{S}^j(t))^c$ is the event that epoch t is *unsuccessful* for node j.

The motivation for this model is that these dynamics capture the queue length evolution under a class of cooperative retransmission MAC protocols discussed in §III. In particular, if a set of nodes $\mathcal{J}_c \subseteq \mathcal{J}$ collide, then the cooperative retransmission protocol ensures that at the end of the collision each of the $|\mathcal{J}_c|$ packets are successfully received at the base station, while the nodes $\mathcal{J} \setminus \mathcal{J}_c$ must remain quiescent until the recovery is completed. The random variables $R_s^j(t), R_u^j(t)$ are not identically distributed: the number of arrivals for successful and unsuccessful epochs depends upon the distribution of the duration of the epochs. The distribution on the duration of an epoch depends upon whether or not node j is participating or not in the epoch.

We employ moment generating functions (mgf) for the random variables $Q^j(t), R_s^j(t), R_u^j(t)$ for each node j and each epoch t:

- $G_{Q^j(t)}(z) = \mathbb{E}\left[z^{Q^j(t)}\right]$ is the mgf for $Q^j(t)$;
- $G_{R_u^j(t)}(z) = \mathbb{E}\left[z^{R_u^j(t)}\right]$ is the mgf for $R_u^j(t)$;
- $G_{R_s^j(t)}(z) = \mathbb{E}\left[z^{R_s^j(t)}\right]$ is the mgf for $R_s^j(t)$.

By applying the mgfs to (1), taking the limit as $t \to \infty$, exploiting the symmetry across nodes, and solving for $G_Q(z)$ we obtain the following result:

Theorem 1: The steady state queue length mgf may be expressed in terms of the arrival mgfs:

$$G_Q(z) = (1-p_a)p_{s|a}\frac{zG_{R_u}(z) - G_{R_s}(z)}{z - p_{s|a}G_{R_s}(z) - (1-p_{s|a})zG_{R_u}(z)} \quad (2)$$

where

- $\mathcal{A}^j(t) = \{Q^j(t) > 0\}$ is the event queue j is non-empty at start of epoch t;
- $p_a = \lim_{t\to\infty}\mathbb{P}(\mathcal{A}^j(t))$ is the probability a queue is non-empty (active);
- $p_{s|a} = \lim_{t\to\infty}\mathbb{P}(\mathcal{S}^j(t)|\mathcal{A}^j(t))$ is the probability epoch t successful given it is active.

An immediate consequence of Theorem 1 is the following corollary relating the probability a typical node is successful to the new packet arrival rate and the average durations of successful and unsuccessful epochs:

Corollary 1: The probability a typical node is successful in a typical epoch (p_s) is given by:

$$p_s = \frac{\lambda\tau_u}{\lambda\tau_u + 1 - \lambda\tau_s}. \quad (3)$$

Here:

- $G'_{R_u}(1) = \mathbb{E}[R_u] = \lambda\mathbb{E}[T_u(t)] = \lambda\tau_u$ is the average number of arrivals in an *unsuccessful* epoch;
- $G'_{R_s}(1) = \mathbb{E}[R_s] = \lambda\mathbb{E}[T_s(t)] = \lambda\tau_s$ is the average number of arrivals in a *successful* epoch.

This corollary is obtained by evaluating (2) at $z = 1$. The number of arrivals in an epoch is simply the average number of arrivals per time slot (λ) times the average number of slots per epoch (τ_u and τ_s). The average epoch durations τ_u and τ_s depends upon the number of nodes (J) and the probabilities p_a and $p_{s|a}$. This equation serves as the basis for computing the stability requirements for the four protocols in §III.

The following theorem gives the average queue length, $\mathbb{E}[Q]$ for a typical node in steady state. The theorem is obtained by differentiating both sides of (2) with respect to z and then taking the limit as $z \to 1$, using $\mathbb{E}[Q] = \lim_{z\to 1} G'_Q(z)$. Taking this limit requires two successive applications of L'Hopital's Rule as well as extensive tedious algebra.

Theorem 2: When the system is stable, the average queue length of a typical node at the start of an epoch in steady state is:

$$\mathbb{E}[Q] = \frac{(1-p_a)p_s}{2} \times \frac{\tau_u(2\lambda+\sigma^2) + \lambda^2(\nu_u - 2\tau_s\tau_u) + \lambda^3(\tau_u\nu_s - \tau_s\nu_u)}{(p_s(1-\lambda\tau_s) - (1-p_s)\lambda\tau_u)^2}. \quad (4)$$

Here:

- $\nu_u = \mathbb{E}[T_u^2(t)]$ is the second moment of the duration (in slots) of an *unsuccessful* epoch;
- $\nu_s = \mathbb{E}[T_s^2(t)]$ is the second moment of the duration (in slots) of a *successful* epoch.

Applying Little's Law ($y_{\text{tot}} \equiv \mathbb{E}[D] = \frac{\mathbb{E}[Q]}{\lambda}$) to Theorem 2 yields the following corollary on the average delay seen by a typical packet at a typical node in steady state:

Corollary 2: When the system is stable, the average delay (queueing plus service) of a typical packet in steady state is

$$y_{\rm tot} = \frac{(1-p_a)p_s}{2\lambda} \times \frac{\tau_u(2\lambda+\sigma^2)+\lambda^2(\nu_u-2\tau_s\tau_u)+\lambda^3(\tau_u\nu_s-\tau_s\nu_u)}{(p_s(1-\lambda\tau_s)-(1-p_s)\lambda\tau_u)^2}. \quad (5)$$

Finally, the following theorem gives the expected aggregate throughput of the system, measured in terms of number of packets per time slot.

Theorem 3: The average throughput of the system (all queues combined) in steady state is:

$$x \equiv \mathbb{E}[X] = \frac{J\mathbb{E}[N]}{\mathbb{E}[T]} = \frac{Jp_s}{\tau_s p_s + \tau_u p_u}. \quad (6)$$

Here:

- $\mathbb{E}[N] = p_s$ is the expected number of packets successfully transmitted by a node in a typical epoch;
- $\mathbb{E}[T] = \tau_s p_s + \tau_u p_u$ is the expected duration (in slots) of a typical epoch.

The theorem follows immediately from the renewal reward theorem for renewal processes, where the first time slot of each epoch is a renewal instant and the reward for each epoch is the number of packets successfully received at the base station at the end of the epoch. To summarize, the queueing dynamics in (1) give rise to the steady state equation for the mgf of the queue length (2), which in turn gives rise to the probability of success (3), the average queue length (4), the average delay (5), and the average aggregate throughput (6).

III. Analysis of four cellular uplink protocols

We now apply the general queueing dynamics from §II to four specific protocols for cellular uplink. In each case we keep all assumptions given at the start of §II, and we further assume each active node randomly elects to attempt transmission in the first time slot of each epoch with probability p_c, which we call the *contention probability*. Inactive nodes (nodes having empty queues at the start of an epoch) do not attempt transmission.

1) *Protocol 1: independent queues.* The J queues do not interact, meaning their transmission attempts do not collide. Thus any node that attempts transmission is automatically successful. This is clearly unrealistic, but the performance of this protocol serves as an upper bound for the other three protocols.
2) *Protocol 2: slotted Aloha.* Time slots are either idle (if no node attempts to transmit), successful (if exactly one node attempts to transmit), or unsuccessful (if multiple nodes attempt to transmit).
3) *Protocol 3: perfect cooperative recovery.* Similar to slotted Aloha, except that a transmission attempt by multiple (say, J_c) nodes triggers a collision resolution epoch of duration J_c time slots (including the colliding time slot), after which all J_c packets are successfully received at the base station.
4) *Protocol 4: imperfect cooperative recovery.* Similar to perfect cooperative recovery, except that the duration of a resolution epoch for J_c colliding nodes has a negative binomial distribution (more specifically, a Pascal distribution). In particular, the epoch ends upon J_c "successful" time slots, where each time slot is independently and identically likely to be a success with probability $1-q$.

We characterize the stability, throughput and delay properties of these four protocols using the general results from §II.

A. Protocol 1: independent queues

Under this protocol all transmission attempts are successful, thus $p_{s|a} = p_c$ (the probability of success given a node is active is the probability an active node attempts to transmit), and so $p_s = p_{s|a}p_a + p_{s|\bar{a}}(1-p_a) = p_c p_a$. Further, each time slot is an epoch of duration 1, hence $\tau_s = \tau_u = \nu_s = \nu_u = 1$. The probability of success equation (3) gives $p_s = \lambda$, or $p_a = \frac{\lambda}{p_c}$. Of course $p_a \in [0,1]$, yielding: $p_a = \min\left\{\frac{\lambda}{p_c}, 1\right\}$. The throughput equation (6) yields

$$x = Jp_s = Jp_c p_a = \min\{J\lambda, Jp_c\}. \quad (7)$$

which states the throughput equals the aggregate arrival rate provided the system is stable ($J\lambda < Jp_c$), else the throughput is saturated at Jp_c. When stable, the delay equation (5) yields:

$$y_{\rm tot} = \frac{(1-\lambda/p_c)p_c}{2\lambda}\frac{2\lambda+\sigma^2-\lambda^2}{(p_c-\lambda)^2} = \frac{2\lambda+\sigma^2-\lambda^2}{2\lambda(p_c-\lambda)} \quad (8)$$

Note delay grows without bound as the arrival rate nears the stability threshold ($\lambda \uparrow p_c$).

B. Protocol 2: slotted Aloha

The elements of a binomial probability mass function are denoted:

$$b_k^n(p) = \mathbb{P}(\mathrm{Bin}(n,p)=k) = \binom{n}{k}p^k(1-p)^{n-k}, \quad (9)$$

for $k = 0,\ldots,n$. Each queue contends with probability p_c when active, so: $p_s = p_{s|a}p_a = p_c p_a(1-p_c p_a)^{J-1} = b_1^J(p_c p_a)/J$. Further, all epochs are again of duration 1, thus: $\tau_s = \tau_u = \nu_s = \nu_u = 1$. The success probability equation (3) becomes $p_s = \lambda$, or $b_1^J(p_c p_a) = J\lambda$. This equation has a real solution for p_a provided $J\lambda < b_1^J(p_c)$:

$$p_a = \begin{cases} \frac{w}{p_c} \text{ where } (b_1^J(w) = J\lambda), & J\lambda < b_1^J(p_c) \\ 1, & \text{else} \end{cases}. \quad (10)$$

Thus the stability requirement is $J\lambda < b_1^J(p_c)$. The throughput equation (6) gives $x = b_1^J(p_c p_a)$, which can be written:

$$x = \begin{cases} J\lambda, & J\lambda < b_1^J(p_c) \\ b_1^J(p_c), & \text{else} \end{cases} = \min\left\{J\lambda, b_1^J(p_c)\right\}. \quad (11)$$

This states the throughput equals the aggregate arrival rate provided the system is stable ($J\lambda < b_1^J(p_c)$), else the throughput is saturated at $b_1^J(p_c)$. The delay equation (5) gives:

$$y_{\rm tot} = \frac{(p_c-w)(1-w)^{J-1}(2\lambda+\sigma^2-\lambda^2)}{2\lambda(p_c(1-w)^{J-1}-\lambda)^2}. \quad (12)$$

C. Protocol 3: perfect cooperative recovery

Each queue contends with probability p_c when active, and all attempted transmissions are (eventually) successful. Thus: $p_s = p_{s|a} p_a = p_c p_a$. The following theorem on the first and second moments of the random durations of the epochs under this protocol is obtained by careful evaluation of the probabilities that the duration of an epoch is j time slots ($\{T = j\}$) in the definitions:

$$\begin{array}{lllll} \tau_u & = & \mathbb{E}[T_u] & = & \sum_{j=1}^{J-1} j\mathbb{P}(T_u = j) \\ \tau_s & = & \mathbb{E}[T_s] & = & \sum_{j=1}^{J} j\mathbb{P}(T_s = j) \\ \nu_u & = & \mathbb{E}[T_u^2] & = & \sum_{j=1}^{J-1} j^2\mathbb{P}(T_u = j) \\ \nu_s & = & \mathbb{E}[T_s^2] & = & \sum_{j=1}^{J} j^2\mathbb{P}(T_s = j) \end{array}. \quad (13)$$

Theorem 4: The first and second moments of the random durations of successful and unsuccessful epochs under the perfect cooperative recovery protocol are:

$$\begin{array}{lll} \tau_u & = & (1-p_c p_a)^{J-1} + (J-1)p_c p_a \\ \tau_s & = & 1 + (J-1)p_c p_a \\ \nu_u & = & (1-p_c p_a)^{J-1} + (J-1)p_c p_a(1+(J-2)p_c p_a) \\ \nu_s & = & 1 + 3(J-1)p_c p_a + (J-1)(J-2)p_c^2 p_a^2 \end{array} \quad (14)$$

Note that $\tau_s > \tau_u$: conditioning on a node participating in a cooperative recovery epoch extends the expected duration of that epoch. With these expected durations we can evaluate the probability of success equation (3) to obtain (after extensive algebra):

$$(1-J\lambda)p_c p_a = \lambda(1-p_c p_a)^J. \quad (15)$$

It is straightforward to establish this equation has a unique solution provided the stability requirement is satisfied:

$$J\lambda < \frac{Jp_c}{Jp_c + b_0^J(p_c)}. \quad (16)$$

It follows that the active probability is:

$$p_a \left\{ \begin{array}{ll} p_a : (1-J\lambda)p_c p_a = \lambda b_0^J(p_c p_a), & J\lambda < \frac{Jp_c}{Jp_c + b_0^J(p_c)} \\ = 1, & \text{else} \end{array} \right. \quad (17)$$

Further, the throughput equation (6) yields:

$$x = \min\left\{ J\lambda, \frac{Jp_c}{Jp_c + b_0^J(p_c)} \right\}. \quad (18)$$

Once again, the throughput is the offered load when stable, and the saturated load when unstable. The delay equation (5) can now be evaluated by substitution of the appropriate expressions for

$$p_a,\ p_{s|a} = p_c,\ p_s = p_c p_a,\ \lambda,\ \sigma^2,\ \tau_u,\ \tau_s,\ \nu_u,\ \nu_s. \quad (19)$$

The resulting expression is unwieldy and lends little insight into the delay and is therefore omitted.

D. Protocol 4: imperfect cooperative recovery

This protocol is identical to the last one, except the durations of the epochs now depends upon the imperfection parameter q. Again we have $p_s = p_{s|a} p_a = p_c p_a$. Computing the first and second moments of the random durations of the epochs given in the following theorem requires careful conditioning on the random number of nodes contending in each epoch:

$$\begin{array}{lllll} \tau_u & = & \mathbb{E}[T_u] & = & \sum_{n=0}^{J-1} \mathbb{E}[T_u|N_u = n]\mathbb{P}(N_u = n) \\ \tau_s & = & \mathbb{E}[T_s] & = & \sum_{n=1}^{J} \mathbb{E}[T_s|N_s = n]\mathbb{P}(N_s = n) \\ \nu_u & = & \mathbb{E}[T_u^2] & = & \sum_{n=0}^{J-1} \mathbb{E}[T_u^2|N_u = n]\mathbb{P}(N_u = n) \\ \nu_s & = & \mathbb{E}[T_s^2] & = & \sum_{n=1}^{J} \mathbb{E}[T_s^2|N_s = n]\mathbb{P}(N_s = n) \end{array}. \quad (20)$$

Theorem 5: The first and second moments of the random durations of successful and unsuccessful epochs under the perfect cooperative recovery protocol are (with $p \equiv p_c p_a$):

$$\begin{array}{lll} \tau_u & = & (1-p)^{J-1} + \frac{(J-1)p}{1-q} - \frac{q}{1-q}(J-1)p(1-p)^{J-2} \\ \tau_s & = & \frac{1}{1-q}((J-1)p+1) - \frac{q}{1-q}(1-p)^{J-1} \\ \nu_u & = & (1-p)^{J-1} + \frac{1}{(1-q)^2}(J-1)p(1+(J-2)p) - \\ & & \frac{q(3-q)}{(1-q)^2}(J-1)p(1-p)^{J-2} + \frac{q}{(1-q)^2}(J-1)p \\ \nu_s & = & (1-p)^{J-1} + \frac{(J-1)p(1+(J-2)p) + 2(J-1)p}{(1-q)^2} + \\ & & \frac{1-(1-p)^{J-1}}{(1-q)^2} + \frac{q\left((J-1)p + 1 - (1-p)^{J-1}\right)}{(1-q)^2} \end{array} \quad (21)$$

Note that the expressions in Theorem 5 recover those in Theorem 4 for $q = 0$, consistent with the fact that the perfect recovery protocol is a special instance of the imperfect recovery protocol. Further note that each of $\tau_u, \tau_s, \nu_u, \nu_s$ diverge to ∞ as $q \uparrow 1$. In words, the mean and variance of the epoch durations grows to infinity as the imperfection in the recovery process becomes increasingly dominant. Applying the expressions in Theorem 5 to the probability of success equation (3) yields, after extensive manipulations, the stability criterion to be:

$$J\lambda < \frac{(1-q)Jp_c}{Jp_c + (1-q)(1-p_c)^J - qJp_c(1-p_c)^{J-1}}. \quad (22)$$

Further, the throughput equation (6) yields:

$$x = \min\left\{ J\lambda, \frac{(1-q)Jp_c}{Jp_c + (1-q)(1-p_c)^J - qJp_c(1-p_c)^{J-1}} \right\}. \quad (23)$$

The expressions in Theorem 5 may be used to compute the delay y_{tot} from (5). The expression is, again, quite lengthy and lends little insight, and is therefore omitted. Clearly the performance under perfect cooperative recovery is superior to both slotted Aloha and performance under imperfect cooperative recovery. Moreover, for sufficiently strong recovery imperfection, slotted Aloha outperforms imperfect cooperative recovery. That is, there exists some q_c such that $q > q_c$ means the costs of cooperative recovery outweigh the benefits.

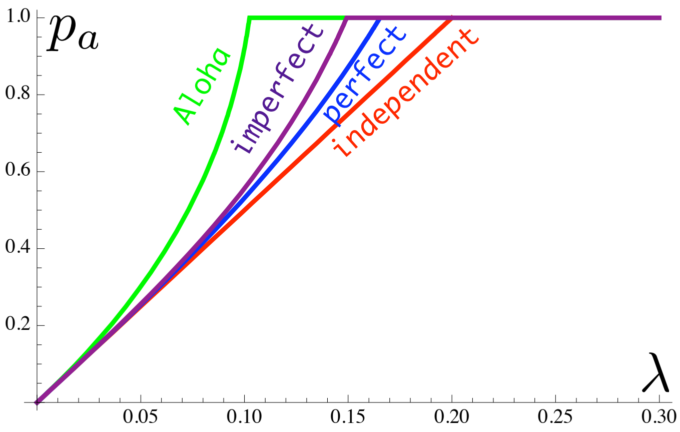

(a) Active probability (p_a) vs. arrival rate (λ).

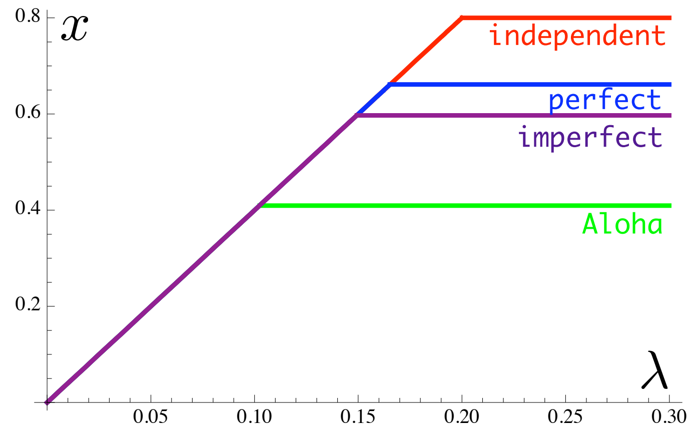

(b) Throughput (x) vs. arrival rate (λ).

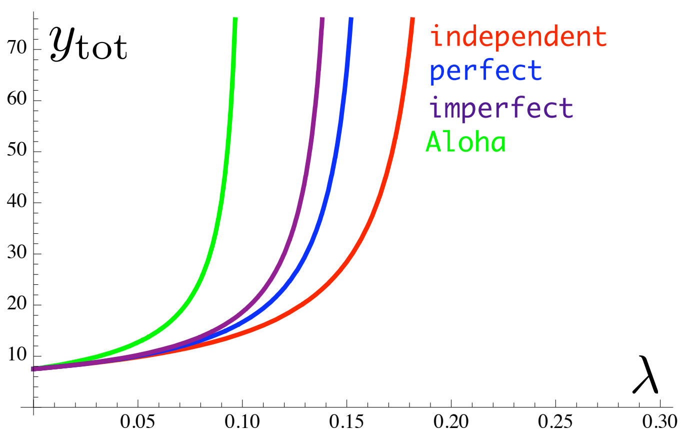

(c) Delay (y_{tot}) vs. arrival rate (λ).

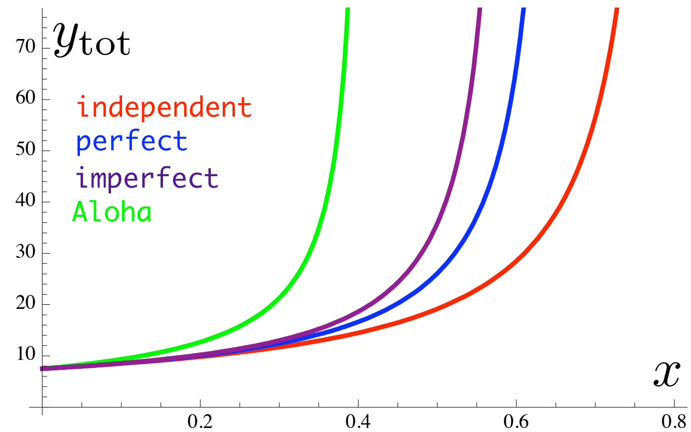

(d) Parametric plot of achievable delay (y_{tot}) and throughput (x).

Fig. 1. Active probability, throughput, and delay (queueing plus service) vs. the arrival rate for the four protocols. For each plot there are $J = 4$ nodes; each active node contends with contention probability $p_c = 0.20$. For the imperfect protocol we use an imperfection parameter of $q = 0.25$.

IV. Numerical results

Figure 1 shows four plots of numerical computations of the expressions in §III. In all plots there are $J = 4$ nodes in the cell, each employing a contention probability $p_c = 0.20$, where the number of arrivals per slot is a Poisson r.v. with parameter λ (hence $\sigma^2 = \lambda$). For the imperfect recovery protocol we used $q = 0.25$. The top left figure shows the active probability p_a versus the arrival rate λ. Note the point where $p_a = 1$ is the critical arrival rate λ_c for stability. Also note the ordering of these critical points:

$$\lambda_c^{\text{Aloha}} < \lambda_c^{\text{Imp}} < \lambda_c^{\text{Per}} < \lambda_c^{\text{Ind}}. \tag{24}$$

The ordering $\lambda_c^{\text{Aloha}} < \lambda_c^{\text{Per}} < \lambda_c^{\text{Ind}}$ will hold for all p_c. For sufficiently high q we would also find $\lambda_c^{\text{Imp}} < \lambda_c^{\text{Aloha}}$. The top right figure shows the throughput x (6) versus the arrival rate λ. As was shown for all four protocols, the throughput equals the aggregate arrival rate ($J\lambda$) provided the system is stable, else the throughput equals the maximum stable aggregate arrival rate ($J\lambda_c$). The bottom left figure shows the total delay (queueing plus service) versus the arrival rate λ. Note that the delay blows up to infinity as $\lambda \uparrow \lambda_c$. Finally, the bottom right plot is a parametric plot of achievable throughput–delay pairs.

V. Conclusion

In §II we considered a general model of queueing dynamics appropriate for a class of cellular uplink protocols. We expressed the mgf for the steady state queue length in terms of the mgf for the number of arrivals in both types of epochs. From here we found expressions for throughput and delay. In §III we specialized these general results to four protocols: independent queues, slotted Aloha, perfect cooperative recovery, and imperfect cooperative recovery. Future work will consider the stability region for imperfect cooperative recovery protocols under heterogeneous arrival rates $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_J)$.

References

[1] M. K. Tsatsanis, R. Zhang, and S. Banerjee, "Network–assisted diversity for random access wireless networks," *IEEE Transactions on Signal Processing*, vol. 48, no. 3, pp. 702–711, March 2000.

[2] H. Yang, A. Petropulu, X. Yang, and T. Camp, "A novel location relay selection scheme for ALLIANCES," *IEEE Transactions on Vehicular Technology*, vol. 57, no. 2, pp. 1272–1284, March 2008.

[3] S. Weber, A. V. Kini, and A. P. Petropulu, "A new approximation for slotted buffered Aloha," in *Proceedings of the Conference on Information Sciences and Systems (CISS)*, Princeton, NJ, March 2008.

[4] ——, "Achievable throughput and service delay for imperfect cooperative retransmission MAC protocols," in *Proceedings of the 3rd ACM Workshop on Performance Monitoring and Measurement of Heterogeneous Wireless and Wired Networks (PM2HW2N)*, Vancouver, BC, October 2008.