Asymptotic analysis of a nonlinear AIMD algorithm

Y. Baryshnikov[1] and E. Coffman[2] and J. Feng[2] and P. Momčilović[3]

[1]*Bell Labs, Lucent Technologies, Murray Hill, NJ 07974.* ymb@research.bell-labs.com
[2]*Electrical Engineering, Columbia University, New York, NY 10027.* {egc,fengjing}@ee.columbia.edu
[3]*EECS, University of Michigan, Ann Arbor, MI 48109.*petar@eecs.umich.edu

The Additive-Increase-Multiplicative Decrease (AIMD) algorithm is an effective technique for controlling competitive access to a shared resource. Let N be the number of users and let $x_i(t)$ be the amount of the resource in possession of the i-th user. The allocations $x_i(t)$ increase linearly until the aggregate demand $\sum_i x_i(t)$ exceeds a given nominal capacity, at which point a user is selected at a random time and its allocation reduced from $x_i(t)$ to $x_i(t)/\gamma$, for some given parameter $\gamma > 1$. In our new, generalized version of AIMD, the choice of users to have their allocations cut is determined by a selection rule whereby the probabilities of selection are proportional to $x_i^\alpha(t)/\sum_j x_j^\alpha$, with α a parameter of the policy. Variations of parameters allows one to adjust fairness under AIMD (as measured for example by the variance of $x_i(t)$) as well as to provide for differentiated service. The primary contribution here is an asymptotic, large-N analysis of the above nonlinear AIMD algorithm within a baseline mathematical model that leads to explicit formulas for the density function governing the allocations $x_i(t)$ in statistical equilibrium. The analysis yields explicit formulas for measures of fairness and several techniques for supplying differentiated service via AIMD.

Keywords: AIMD analysis, congestion avoidance algorithms, fair resource allocation, differentiated service

1 Introduction

A number $N > 1$ of users share a limited resource such as transmission bandwidth, storage, or processing rate. Sharing is competitive, each user requesting more of the resource than any fair allocation of it would provide. Competition creates increasing allocations, with the rate of increase fixed at some constant. A rate of 1 is a convenient normalization to be used hereafter. Competition has to be controlled so that total demand does not become too large. When total demand exceeds the nominal capacity, at time t say, a user is selected according to some policy and its allocation $x(t)$ is reduced to $x(t)/\gamma$, where $\gamma \in (1, \infty)$ is a parameter of the policy.

In a simple example of such AIMD (Additive Increase, Multiplicative Decrease) congestion control policies, a fixed capacity B limits the total demand; if at some time t, total demand reaches level B, a user is selected at random, say the i-th is chosen, and has its allocation $x_i(t)$ reduced to $x_i(t)/\gamma$, so total demand $X_N(t) = \sum_{1\leqslant i\leqslant N} x_i(t)$ is reduced to $B - \frac{\gamma-1}{\gamma}x_i(t)$. A sample function for the resource allocated to a user resembles a simple sawtooth waveform in which linear increases in allocation are punctuated by abrupt decreases by a factor of γ. The total demand is also a sawtooth waveform, but the rate of increase is N and the peaks occur at level B.

Although investigations of specific applications of AIMD policies are not within the scope of this paper, it must be noted that, by far, the most studied application is the Transmission Control Protocol (TCP) of the Internet (Jacobson (1988)), where the discrete nature of packet transmission is replaced by the continuous (fluid flow) model in the formulation above. Users are *windows* which bound the number of packets in transit in the network between acknowledgements.

The literature on modeling of TCP is extensive. Models range from the macroscopic (Fredj et al. (2001)) to the microscopic (Gilbert and Karloff (2003)) with various models focusing on different aspects of the protocol. We refer the reader to Ott et al. (1996); Misra et al. (1999); Padhye et al. (2000); Lakshman et al. (2000); Low (2000); Altman et al. (2000); Yang and Lam (2000); Hollot et al. (2001); Dumas et al. (2002); Guillemin et al. (2004); Tinnakornsrisuphap and La (2004) and references therein. The analysis in this paper is in the spirit of Adjih et al. (2001) and Baccelli et al. (2002) (see also McDonald and Reynier (to appear)), i.e., we consider a mean-field analysis of many users interacting through a single resource. However, in contrast with the earlier papers, our focus is not on modeling TCP, but on studying, in a general setting, the effects that variation of parameters have on the resource partition produced by the AIMD algorithm. And in particular, we examine a new, nonlinear version of the algorithm.

Interestingly, in addition to the last three papers mentioned above in connection with TCP (see also Dumas et al. (2002)), results of similar form apply to the asymptotic analysis of Ethernet's exponential back-off protocol (Aldous (1987)), models in the mathematics of finance (Carmona et al. (1997)), probabilistic counting in data base applications (Flajolet and Martin (1985)), and polling systems (Litvak and van Zwet (2004)). This suggests an underlying common structure not yet fully understood. See, however, recent work (Robert (to appear)) in which the connective tissue of some of these problems is shown to lie in the asymptotic analysis of recursive splitting (or tree) algorithms.

Our study of AIMD, unfettered by application-specific issues, concentrates on baseline models applicable to a broad class of service systems providing a shared resource. In the design of control algorithms, this type of analysis represents a critical first step in the process of deciding whether or not to adopt an AIMD policy. It can be expected to yield qualitative insights that will suggest the general behavior one should anticipate in applications.

The analysis here centers on issues of fairness, differentiated service, and the effects of the parameter γ that determines the decreases in allocations. The fairness question relates to the shares received by different users, and usually implies nearly equal shares by some statistical measure. This problem is addressed in part by a simple boundary protocol whereby the larger shares tend to be selected for reduction. An exogenous marking process randomizes size-cutting epochs. In our case the 'hard' boundary B of our earlier example is made a 'soft' boundary as follows: Whenever the total demand $X_N(t)$ exceeds B, an inhomogeneous Poisson process is activated with a rate parameter proportional to the overflow beyond B, i.e., a rate $\lambda(X_N(t) - B)^+$, where $\lambda > 0$ is a given constant. At each epoch of the Poisson process, a user is chosen at random; in the new, nonlinear version of AIMD considered here, a user is chosen with a probability equal to the α *fraction* of the resource that it currently holds, i.e., the fraction $x_i^\alpha(t)/\sum_{j=1}^N x_j^\alpha(t)$ for the i-th user. The user selected has its share reduced by a factor γ. The Poisson process becomes dormant as soon as the total demand falls below B again. When $\alpha = 0$, the earlier example applies; users are chosen at random, each with probability $1/N$. When $\alpha = 1$ the probability of being selected reduces to the linear special case $x_i(t)/X_N(t)$, i.e., the selection probability increases linearly with size. Two immediate properties are noteworthy: First, the process $X_N(t)$ has the property of an Ornstein-Uhlenbeck process whereby the "restoring force" acting toward the boundary in the overflow region increases with the distance from the boundary; and second, user selection, by its preference for large shares, creates a distribution of share size having a smaller variance than that achieved under random selection. The preference for large sizes increases with α.

To this point, we have assumed symmetric, undifferentiated sharing under AIMD. That is, all users have been assumed to have the same parameters, e.g., unit growth rate and coefficient γ. This case is covered in the next section along with a discussion of AIMD behavior as it varies with parameters. The analytical approach is similar to earlier studies, but the results are explicit and new, with novel applications. Specifically, the parametrization of the dependence of the size-reduction criterion (i.e., the dependence of the parameter α) is new.

In section 3, we examine cases where users are not uniform in their congestion control; randomizing reductions, varying rates, and varying the size-dependence parameter α are all studied within the mathematical model. The last of these is proposed as a new, implementable mechanism for service differentiation. Section 4 concludes the paper with a brief discussion of AIMD in the TCP context by arguing the feasibility of implementing quadratic selection ($\alpha = 2$).

2 Symmetric sharing

2.1 Allocation-size density

Let $\{T_n\}$ be the sequence of times when signals are generated by the marking process (the Poisson process described in the previous section), and let $X_N^{(\alpha)}(t) := \sum_{1\leqslant i\leqslant N} x_i^\alpha(t)$. Then for any $i,\ 1 \leqslant i \leqslant N$, the following serves as a formal reference for the earlier definitions

$$x_i(t) = x_i(T_n) + t - T_n,\ \ t \in [T_n, T_{n+1}),$$
$$x_i(T_{n+1}) = x_i(T_{n+1}^-)\left(1 - (1-\gamma^{-1})1_{\{i=S\}}\right),$$

where $1_{\{\cdot\}}$ is the indicator function, T^- denotes the left limiting value at T, and, for U uniformly distributed on $[0, 1]$,

$$S = \min\left\{k :\ \sum_{m=1}^{k} x_m^\alpha(T_{n+1}^-) \geqslant U X_N^{(\alpha)}(T_{n+1}^-)\right\}.$$

Times between signals consist of two components

$$T_{n+1} - T_n = \frac{1}{N}(B - X_N(T_n))^+ + \tau(X_N(T_n));$$

the first term is the time needed for the intensity of the marking process to become non-zero, and the second is a random interval between the time that the intensity of the marking process becomes non-zero and the time when a signal is generated. The random variable $\tau(\cdot)$ is defined by the tail probability

$$\mathbb{P}[\tau(x) > t] = e^{-\lambda\left(\frac{t^2}{2}+(x-B)^+t\right)},\quad t > 0,$$

of the interarrival time in a non-stationary Poisson process (see, e.g., (Cinlar, 1975, Sec. 4.7)). The parameter of the random variable $\tau(\cdot)$ determines the initial intensity of the Poisson marking process. In particular, for all values of this parameter below B, the initial intensity is 0. We focus on the case when the number of users N is large. In order to allow for finite user allocations, the total capacity needs to be scaled proportionally with N, i.e., $B = Nb$ for some $b > 0$. We relate the number of users with allocations not exceeding some value s at time $t + \Delta t$ to the number with allocations not exceeding s at time t. The dynamics of the process yields, as $\Delta t \downarrow 0$,

$$\mathbb{E}\left[\sum_{i=1}^{N} 1_{\{x_i(t+\Delta t)\leqslant s\}}\right] = \mathbb{E}\left[\sum_{i=1}^{N} 1_{\{x_i(t)\leqslant s-\Delta t\}}\right] + \lambda\Delta t\, \mathbb{E}\left[(X_N(t) - Nb)^+ \cdot \sum_{i=1}^{N} 1_{\{s\leqslant x_i(t)\leqslant \gamma s\}} \frac{x_i^\alpha(t)}{X_N^{(\alpha)}(t)}\right] + 0(\Delta t). \tag{1}$$

The first term on the right-hand side of (1) indicates that users with allocations no more than $s - \Delta t$ at time t have their allocations limited by s at time $t + \Delta t$. This is a simple consequence of linear growth at unit rate. The second term represents the event that a user has its allocation reduced below s within time interval Δt; in that case, prior to the cut, the allocation cannot be larger than γs. If we divide (1) by N, let $N \to \infty$, and introduce the limiting distribution

$$G(s, t) = \lim_{N\to\infty} \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N} 1_{\{x_i(t)\leqslant s\}}\right],$$

then we obtain

$$\frac{G(s,t+\Delta t)}{\Delta t} = \frac{G(s-\Delta t,t)}{\Delta t} + \lambda \frac{(\overline{x}(t)-b)^+}{\overline{x^\alpha}(t)}(1+o(1)) \int_s^{\gamma s} x^\alpha \, d_x G(x,t),$$

as $\Delta t \downarrow 0$, where $\overline{x^\alpha}(t) = \lim_{N\to\infty} X_N^{(\alpha)}(t)/N$. Furthermore, in the limit as $\Delta t \downarrow 0$, we obtain the following partial differential equation

$$\frac{\partial G(s,t)}{\partial t} = -\frac{\partial G(s,t)}{\partial s} + \lambda \frac{(\overline{x}(t)-b)^+}{\overline{x^\alpha}(t)} \int_s^{\gamma s} x^\alpha \, d_x G(x,t). \tag{2}$$

A rigorous proof of convergence and the existence of the limit is beyond the scope of this paper. Assuming these, in the limit as $t \to \infty$, the partial derivative $\partial G/\partial t$ vanishes and $\partial G/\partial s$ tends to a density $g(s)$ described by

Lemma 1. *The density $g(s)$ satisfies*

$$g(s) = \lambda \frac{(\overline{x}-b)^+}{\overline{x^\alpha}} \int_s^{\gamma s} x^\alpha g(x)\, dx, \quad s \geqslant 0, \tag{3}$$

where $\overline{x^\alpha} := \int_0^\infty x^\alpha g(x)\, dx$.

In order to obtain a closed form expression for $g(\cdot)$ we establish the following auxiliary result.

Lemma 2. *Let $\beta = \gamma^{\alpha+1}$. The function*

$$h(s) = \sum_{k=0}^{\infty} \frac{\beta^k}{\prod_{i=1}^{k}(1-\beta^i)} e^{-c\frac{\beta^k}{\alpha+1}s^{\alpha+1}}, \quad s \geqslant 0, \tag{4}$$

is a positive integrable solution (unique up to a constant factor in the class of integrable functions rapidly decreasing at infinity) of the integral equation

$$h(s) = c \int_s^{\gamma s} x^\alpha h(x)\, dx, \quad c > 0. \tag{5}$$

Proof. Equation (5) implies a system of linear relations for the moments of h, from which the uniqueness follows via standard arguments. That (4) solves (5) can be verified immediately. ☐

Denote the space of solutions to (5) as $\mathcal{S}(c,\alpha,\gamma),\ c>0,\ \alpha>-1,\ \gamma>1$. It turns out that the group of log-linear substitutions

$$g_{k,\rho}(h)(x) := h(kx^\rho)$$

permutes the spaces S as follows:

Lemma 3. *If h belongs to $\mathcal{S}(c,\alpha,\gamma)$ (i.e. h solves (5)), then*

$$g_{k,\rho}(h)(x) = h(kx^\rho) \in \mathcal{S}(\rho k^{\alpha+1}c, (\alpha+1)\rho - 1, \gamma^{1/\rho})$$

This group action on the solutions of (5) accounts for most of their diversity.

The following theorem characterizes, based on Lemma 1 and Lemma 2, the density $g(\cdot)$. Let $\Gamma(\cdot)$ denote the Gamma function and, for notational simplicity, for given γ, let

$$q(x) := \sum_{k=0}^{\infty} \frac{\gamma^{xk}}{\prod_{i=1}^{k}(1-\gamma^{xi})}.$$

Theorem 1. *The density $g(s)$ is given by*

$$g(s) = \mu \sum_{k=0}^{\infty} \frac{\gamma^{(\alpha+1)k}}{\prod_{i=1}^{k}(1-\gamma^{(\alpha+1)i})} e^{-\frac{\gamma^{(\alpha+1)k}}{\alpha+1}(\xi s)^{\alpha+1}}, \quad s \geqslant 0, \tag{6}$$

where

$$\mu = \frac{(\alpha+1)^{\frac{\alpha}{\alpha+1}}\xi}{\Gamma\left(\frac{1}{\alpha+1}\right) q(\alpha)} \tag{7}$$

and ξ is given by

$$\xi = \frac{\lambda b}{2} \frac{q(\alpha)}{q(0)} \frac{\Gamma(\frac{1}{\alpha+1})}{(\alpha+1)^{\frac{\alpha}{\alpha+1}}} \left[\sqrt{1 + \frac{4(\alpha+1)}{\lambda b^2} \frac{q(\alpha-1)q(0)}{q^2(\alpha)} \frac{\Gamma(\frac{2}{\alpha+1})}{\Gamma^2(\frac{1}{\alpha+1})}} - 1 \right]. \tag{8}$$

Proof. Integral equations (3) and (5) are of the same form and, thus, $g(s)$ can be represented as in (6) with ξ and μ being positive constants. To ensure that $g(s)$ is a probability density, i.e. integrates to 1, the constant μ needs to satisfy (recall that $\beta = \gamma^{\alpha+1}$)

$$\begin{aligned}
\mu^{-1} &= \sum_{k=0}^{\infty} \frac{\beta^k}{\prod_{i=1}^{k}(1-\beta^i)} \int_0^{\infty} e^{-\frac{\beta^k}{\alpha+1}(\xi s)^{\alpha+1}}\, ds \\
&= \frac{1}{\alpha+1} \sum_{k=0}^{\infty} \frac{\beta^k}{\prod_{i=1}^{k}(1-\beta^i)} \left(\frac{\alpha+1}{\beta^k}\right)^{\frac{1}{\alpha+1}} \Gamma\left(\frac{1}{\alpha+1}\right) \frac{1}{\xi} = \frac{\Gamma\left(\frac{1}{\alpha+1}\right) q(\alpha)}{(\alpha+1)^{\frac{\alpha}{\alpha+1}}\xi}
\end{aligned}$$

which is the reciprocal of (7), as desired; the second equality follows from integration by parts. The remaining free parameter ξ is selected in such a way that (3) holds (see also (5)). To this end, the rth moment of $g(\cdot)$ is given by

$$\begin{aligned}
\int_0^{\infty} s^r g(s)\, ds &= \mu \sum_{k=0}^{\infty} \frac{\beta^k}{\prod_{i=1}^{k}(1-\beta^i)} \int_0^{\infty} s^r e^{-\frac{\beta^k}{\alpha+1}(\xi s)^{\alpha+1}}\, ds \\
&= \frac{\mu}{\alpha+1}(\alpha+1)^{\frac{r+1}{\alpha+1}} \frac{q(\alpha-r)}{\xi^{r+1}} \Gamma\left(\frac{r+1}{\alpha+1}\right) = \frac{(\alpha+1)^{\frac{r}{\alpha+1}}}{\xi^r} \frac{q(\alpha-r)}{q(\alpha)} \frac{\Gamma\left(\frac{r+1}{\alpha+1}\right)}{\Gamma\left(\frac{1}{\alpha+1}\right)},
\end{aligned} \tag{9}$$

after substitution for μ from (7). Hence, setting $r = 1$ to evaluate $\overline{x}$, Lemmas 1 and 2 show that ξ must satisfy

$$\lambda \left(\frac{(\alpha+1)^{\frac{1}{\alpha+1}}}{\xi} \frac{q(\alpha-1)}{q(\alpha)} \frac{\Gamma\left(\frac{2}{\alpha+1}\right)}{\Gamma\left(\frac{1}{\alpha+1}\right)} - b \right) = \xi\, (\alpha+1)^{\frac{\alpha}{\alpha+1}} \frac{q(0)}{q(\alpha)} \frac{1}{\Gamma\left(\frac{1}{\alpha+1}\right)},$$

or equivalently, ξ is the positive solution to the quadratic equation

$$\xi^2 + \xi\lambda b \frac{q(\alpha)}{q(0)} \frac{\Gamma(\frac{1}{\alpha+1})}{(\alpha+1)^{\frac{\alpha}{\alpha+1}}} - \lambda \frac{q(\alpha-1)}{q(0)} \frac{\Gamma(\frac{2}{\alpha+1})}{(\alpha+1)^{\frac{\alpha-1}{\alpha+1}}} = 0. \tag{10}$$

The solution of (10) is given by (8) as desired. □

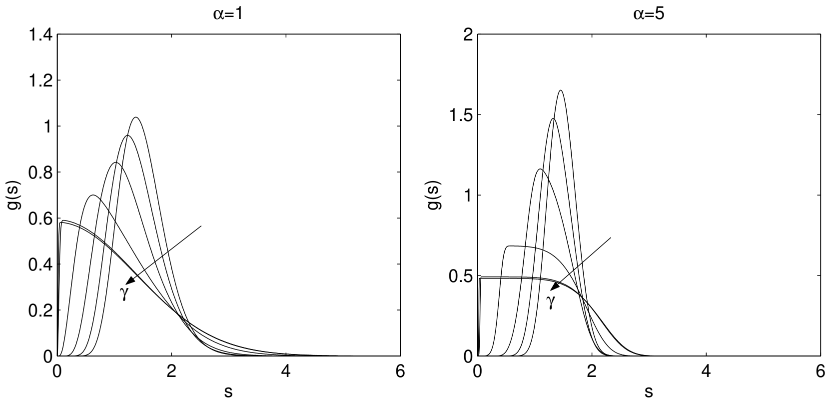

Fig. 1: Density $g(s)$ with $\gamma = 4/3$, 3/2, 2, 5, 50, and ∞ as a parameter. The curves corresponding to $\gamma = \infty$ are plots of (13). Parameter α is set to 1 and 5 for the left and right plot, respectively.

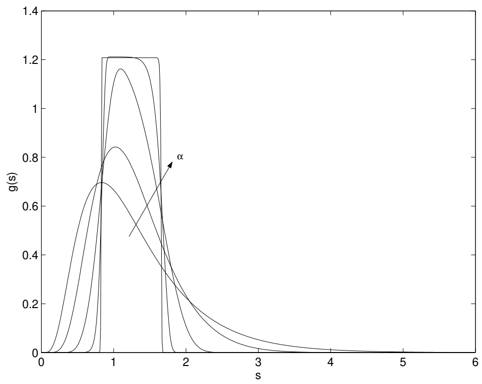

Fig. 2: Density $g(s)$ for $\gamma = 2$ and $\alpha = 0, 1, 5, 20, 200$.

The density $g(s)$ is graphed in Figures 1 and 2 for a variety of values of the parameters γ and α. The general shapes of these densities resemble those of Adjih et al. (2001); Baccelli et al. (2002); Guillemin et al. (2004). However, the details of the analysis are different, and in the model here, the algorithm allows for nonlinear user selection.

The value at the origin, $g(0) = 0$, is not immediately obvious from the formula, but is in fact easy to prove; it comes down to proving that $q(\alpha+1) = 0$. To verify this last equality, write $q(\alpha+1) = \lim_{n\to\infty} q^{(n)}(\alpha+1)$, where

$$q^{(n)}(\alpha+1) = \sum_{k=0}^{n} \frac{\gamma^{(\alpha+1)k}}{\prod_{i=1}^{k}(1-\gamma^{(\alpha+1)i})}, \quad n \geqslant 1.$$

An easy induction establishes that for $n \geqslant 1$

$$q^{(n)}(\alpha+1) = \left(\prod_{i=1}^{n}(1-\gamma^{(\alpha+1)i})\right)^{-1},$$

and since $\gamma > 1$, it follows that $q(\alpha+1) = 0$, as claimed.

Observe from (10) that, in the limit as $\lambda \to \infty$, one has

$$\xi \to \frac{(\alpha+1)^{\frac{1}{\alpha+1}}}{b} \frac{\Gamma(\frac{2}{\alpha+1})}{\Gamma(\frac{1}{\alpha+1})} \frac{q(\alpha-1)}{q(\alpha)}$$

and, therefore, from (9), with $r = 1$, we have that $\overline{x} \to b$, as expected.

2.2 Specific policies

Specializing the result of Theorem 1 provides results for a number of potentially implementable policies.

Random selection When users are chosen uniformly at random ($\alpha = 0$), then $\mu = \xi/q(0)$ and

$$\xi = \frac{\lambda b}{2} \left[\sqrt{1 + \frac{4}{\lambda b^2} \frac{q(-1)}{q(0)}} - 1 \right].$$

Linear selection In the case $\alpha = 1$, the selection probability is proportional to size, and the coefficients are $\mu = \sqrt{2/\pi}\, \xi/q(1)$ and

$$\xi = \frac{\lambda b}{2} \sqrt{\frac{\pi}{2}} \frac{q(1)}{q(0)} \left(\sqrt{1 + \frac{8}{\lambda b^2 \pi} \frac{q^2(0)}{q^2(1)}} - 1 \right).$$

This policy with size bisection ($\gamma = 2$) is the one most often seen in the literature.

Largest-allocation selection As suggested by Figure 2, $g(s)$ simplifies to a uniform distribution for this selection rule, which is obtained in the limit as $\alpha \to \infty$.

Corollary 1. *In the limit $\alpha \to \infty$, the density $g(s)$ is given by*

$$g(s) = \mu(1_{\{s\mu > 1/(\gamma-1)\}} - 1_{\{s\mu > \gamma/(\gamma-1)\}}), \tag{11}$$

i.e., it is uniform on $(\frac{1}{\mu(\gamma-1)}, \frac{\gamma}{\mu(\gamma-1)}]$, where

$$\mu = \frac{\lambda b}{2} \left[\sqrt{1 + \frac{2}{\lambda b^2} \frac{\gamma+1}{\gamma-1}} - 1 \right]. \tag{12}$$

Proof. We need the estimates $\Gamma(x) = x^{-1} + o(x^{-1})$, as $x \downarrow 0$, and $q(\alpha - i) = 1 - \gamma^{-i-1} + o(\gamma^{-i-1})$, as $\alpha \to \infty$, to obtain

$$\xi = \frac{\lambda b}{2} \frac{\gamma-1}{\gamma} \left[\sqrt{1 + \frac{2}{\lambda b^2} \frac{\gamma+1}{\gamma-1}} - 1 \right] (1 + o(1)),$$

as $\alpha \to \infty$. Hence, the formula for μ given in (12). According to Theorem 1 the density $g(s)$ can be represented as an infinite sum. However, in the limit as $\alpha \to \infty$ only the first two terms have nonzero limits and the expression reduces to (11). □

The qualitative behavior corresponding to this $\alpha \to \infty$ result is easily described. For large N, the users can be thought of as being organized into a cyclic chain. Following the user ahead of it(by a time about one N-th the cycle time), each user increases its allocation from about $2\overline{x}/(\gamma+1)$ to about $2\gamma\overline{x}/(\gamma+1)$, then has it cut back to about $2\overline{x}/(\gamma+1)$, and repeats the process cyclically. Thus, in the limit, $g(s)$ converges to a uniform distribution on $(\frac{2\overline{x}}{\gamma+1}, \frac{2\gamma\overline{x}}{\gamma+1}]$. In this case $\overline{x} = \frac{1}{2\mu} \frac{\gamma+1}{\gamma-1}$.

Allocation reset In the limit $\gamma \to \infty$ we arrive at a policy in which allocations selected for reduction are reset to (are made to restart from) zero. In this limit, only the $k = 0$ term survives in (6) and in the sum that defines $q(x)$. In this limit then, $q(x) = 1$, $0 \leqslant x \leqslant \alpha$, and

$$g(s) = \frac{(\alpha+1)^{\frac{\alpha}{\alpha+1}}\xi}{\Gamma(\frac{1}{\alpha+1})} e^{-\frac{(\xi s)^{\alpha+1}}{\alpha+1}}, \quad s > 0, \tag{13}$$

where

$$\xi = \frac{\lambda b}{2} \frac{\Gamma(\frac{1}{\alpha+1})}{(\alpha+1)^{\frac{\alpha}{\alpha+1}}} \left[\sqrt{1 + \frac{4(\alpha+1)}{\lambda b^2} \frac{\Gamma(\frac{2}{\alpha+1})}{\Gamma^2(\frac{1}{\alpha+1})}} - 1 \right].$$

This formula for $g(s)$ is plotted in Figure 1 for the limit $\gamma \to \infty$.

Jitter limit If $\gamma = 1 + \epsilon, \epsilon \to 0$, with fixed c, α, the solutions of (5), properly rescaled, tend to a Gaussian curve. This is enough to prove for a fixed α (by virtue of lemma 3), for example for $\alpha = 0$. In this case, it is easy to prove that the deviation $Y(t)$ of the size of a given window from its mean (which is of order ϵ^{-1}), when rescaled as $\epsilon^{1/2}Y(t/\epsilon)$, converges to a non-degenerate Ornstein-Uhlenbeck process.

2.3 Variability and fairness

Theorem 1 allows one to estimate the variability of allocations observed by individual users. The coefficient of variation, the ratio of the standard deviation to the mean, is a conventional measure, and by (9) it is given by

$$\rho = \sqrt{\frac{\overline{x^2}}{\overline{x}^2} - 1} = \sqrt{\frac{q(\alpha-2)q(\alpha)}{q^2(\alpha-1)} \frac{\Gamma(\frac{3}{\alpha+1})\Gamma(\frac{1}{\alpha+1})}{\Gamma^2(\frac{2}{\alpha+1})} - 1}, \tag{14}$$

which we note is independent of the parameters λ and b. As expected, ρ is monotone increasing in γ and decreasing in α.

A lower bound on the coefficient of variation is as follows.

Theorem 2. *The coefficient of variation satisfies*

$$\rho \geqslant \frac{1}{\sqrt{3}} \frac{\gamma - 1}{\gamma + 1}. \tag{15}$$

The bound is achieved by largest-allocation selection, i.e., in the limit as $\alpha \to \infty$.

Proof. For a given γ and a given average allocation $\overline{x}$, a lower bound to the coefficient of variation is given by the (deterministic) periodic saw-tooth

$$x(t) = t - 2\overline{x}\frac{\gamma-1}{\gamma+1} \sum_{i=1}^{\infty} 1_{\{t \geqslant 2\overline{x} i \frac{\gamma-1}{\gamma+1}\}} + \frac{2\overline{x}}{\gamma+1},$$

which rises at unit rate from $\frac{2\overline{x}}{\gamma+1}$ to $\frac{2\gamma\overline{x}}{\gamma+1}$ in every time interval

$$\left[\frac{2\overline{x} i(\gamma-1)}{\gamma+1}, \frac{2\overline{x}(i+1)(\gamma-1)}{\gamma+1} \right), \quad i \geqslant 0.$$

A routine calculation of $\overline{x^2}/\overline{x}^2$ for this function shows that (15) holds. To verify that, in the limit $\alpha \to \infty$, this lower bound is achieved, introduce into (14) the earlier estimates for $\Gamma(x)$ and $q(\alpha - i)$ to obtain

$$\lim_{\alpha \to \infty} \rho = \sqrt{\frac{4}{3} \frac{(\gamma-1)(\gamma^3-1)}{(\gamma^2-1)^2} - 1} = \frac{1}{\sqrt{3}} \frac{\gamma-1}{\gamma+1}. \qquad \square$$

As examples, note that, for $\alpha = \infty$, size bisection gives $\rho = \frac{1}{3\sqrt{3}}$, and size reset gives $\rho = \frac{1}{\sqrt{3}}$.

3 Service differentiation

Proper functioning of a distributed application of the AIMD algorithm is based on the willingness of users to participate. This motivates the extensions below to non-uniform users and varying grades of service. Let the set of all users be partitioned into K groups such that users within a given group behave in the same fashion. The size of the i-th group is given by $\theta_i N$ with $\theta_i > 0$ for all i. (Hereafter, i and j are reserved for indexing groups, and are attached as needed to the parameters of the preceding section.) In order to quantify the advantage gained by different groups we define $\chi_{ij}(r)$ as the ratio of the rth moments of allocations of users in groups i and j, i.e.,

$$\chi_{ij}(r) := \frac{\int_0^\infty x^r g_i(x)\,dx}{\int_0^\infty x^r g_j(x)\,dx}.$$

3.1 Varying randomization

A selected user from the i-th group reduces its size (by the factor γ) only with probability p_i. Following the analysis in the preceding section, it is straightforward to conclude that densities $g_i(\cdot)$ of allocations for users in group i satisfy the following equation

$$g_i(s) = \lambda p_i \frac{(\overline{x} - b)^+}{\overline{x^\alpha}} \int_s^{\gamma s} x^\alpha g_i(x)\,dx, \tag{16}$$

where $\overline{x^\alpha}$ is the average over all users, i.e., $\overline{x^\alpha} = \sum_{i=1}^K \theta_i \int_0^\infty x^\alpha g_i(x)\,dx$. Based on (9) and (16) it is easy to see that $\chi_{ij}(r) = (\xi_j/\xi_i)^r = (p_j/p_i)^{r/(\alpha+1)}$. We note in passing that this formula is related to the well-known expression ($\sqrt{p}$-formula (Floyd (1991))) for the throughput of TCP connections. Specifically, for linear selection ($\alpha = 1$) and the throughput measure ($r = 1$), one has $\chi_{ij}(1) = \sqrt{p_j/p_i}$.

3.2 Varying growth rates

For varying rates of increase, let the rate be κ_i for users in the i-th group (all other parameters are group independent). Then $\chi_{ij} = (\kappa_i/\kappa_j)^{r/(\alpha+1)}$, since the functions $g_i(\cdot)$ satisfy

$$g_i(s) = \lambda \kappa_i^{-1} \frac{(\overline{x} - b)^+}{\overline{x^\alpha}} \int_s^{\gamma s} x^\alpha g_i(x)\,dx$$

3.3 Varying size preference

Finally, consider the case in which different groups of users employ different parameters α (all other parameters are the same for all groups). In this case a user is chosen to reduce its rate with a probability proportional to its α_i allocation, where i is the index of the group to which the user belongs. From (9) it follows that

$$\chi_{ij}(r) := \frac{\int_0^\infty x^r g_i(x)\,dx}{\int_0^\infty x^r g_j(x)\,dx} = \left(\frac{\xi_j}{\xi_i}\right)^r,$$

where the dependence on i is restricted to $\xi_i = \eta^{\alpha_i}$ with η the same for all groups.

4 Concluding remarks

Our current research has turned to applications, initially the TCP application, within realistic simulation models. The Random Early Detection (RED) protocol (Floyd and Jacobson (1995)) provides a framework for testing nonlinear AIMD in practice. The packet flows of N source-destination pairs share the bandwidth supplied by the network. According to RED, when the bound B is exceeded, a packet is marked at random by a router; when received a marked packet causes a reduction in size of the corresponding window by the factor γ. Marking packets uniformly at random means that the probability of marking a packet from a given window, and hence reducing the size of that window, becomes proportional to the window's size; this is the linear case $\alpha = 1$. This mechanism can be extended to $\alpha = 2$. A source does not change its window size in response to at most one marked packet; only when it receives two or more marked packets of a window does it make a reduction in window size. The result is quadratic selection whereby the probability of window reduction increases as the square of window size.

References

C. Adjih, P. Jacquet, and N. Vvedenskaya. Performance evaluation of a single queue under multi-user TCP/IP connections. Technical Report RR-4141, INRIA, March 2001.

D. Aldous. Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels. *J. Information Theory*, 33:219–223, 1987.

E. Altman, K. Avrachenkov, and C. Barakat. A stochastic model of TCP/IP with stationary random losses. In *Proc. ACM SIGCOMM 2000*, Stockholm, Sweden, August 2000.

F. Baccelli, D. McDonald, and J. Reynier. A mean-field model for multiple TCP commenctions through a buffer implementing RED. In *Proc. Performance*, Rome, Italy, September 2002.

P. Carmona, F. Petit, and M. Yor. On the distribution and asymptotic results for exponential functionals of Lévy processes. In *Exponential Functionals and Principal Values Related to Brownian Motion (Rev. Mat. Iberoam)*, pages 73–130. Madrid, 1997.

E. Cinlar. *Introduction to Stochastic Processes*. Prentice-Hall, 1975.

V. Dumas, F. Guillemin, and P. Robert. A markovian analysis of additive-increase multiplicative-decrease algorithms. *Ann. Appl. Probab.*, 34:85–111, 2002.

P. Flajolet and N. Martin. Probabilistic counting for data base applications. *J. Computer Sys.*, 31:182–209, 1985.

S. Floyd. Connections with multiple congested gateways in packet-switched networks. Part 1: One-way traffic. *Computer Comm. Review*, 21(5):30–47, 1991.

S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, 1995.

S. B. Fredj, T. Bonald, A. Proutiere, G. Regnie, and J. Roberts. Statistical bandwidth sharing: A study of congestion at flow level. In *Proc. ACM SIGCOMM*, San Diego, CA, August 2001.

A. Gilbert and H. Karloff. On the fractal behavior of TCP. In *Proc. ACM STOC*, San Diego, CA, June 2003.

F. Guillemin, P. Robert, and B. Zwart. AIMD algorithms and exponential functions. *Ann. Appl. Probab.*, 14(1):90–117, 2004.

C. Hollot, V. Misra, D. Towsley, and W. Gong. A control theoretic analysis of RED. In *Proc. IEEE INFOCOM*, Anchorage, AK, April 2001.

V. Jacobson. Congestion avoidance and control. In *Proc. ACM SIGCOMM*, pages 314–329, Stanford, CA, August 1988.

T. Lakshman, U. Madhow, and B. Suter. TCP/IP performance with random loss and bidirectional congestion. *IEEE/ACM Transactions on Networking*, 8(5):541–555, October 2000.

N. Litvak and W. van Zwet. On the minimal travel time needed to collect n items on a circle. *Ann. Appl. Probab.*, 14:881–902, 2004.

S. Low. A duality model of TCP and queue management algorithms. In *ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, Monterey, CA, September 2000.

D. McDonald and J. Reynier. Mean filed convergence of a rate of multiple TCP connections through a buffer implementing RED. *Ann. Appl. Probab.*, to appear.

V. Misra, W. Gong, and D. Towsley. Stochastic differential equation modeling and analysis of TCP window size behavior. In *Proc. Performance'99*, Istanbul, Turkey, September 1999.

T. Ott, J. Kemperman, and M. Mathis. The stationary behavior of ideal TCP congestion avoidance. , August 1996.

J. Padhye, V. Firoiu, and D. Towsley. Modeling TCP Reno performance: A simple model and its empirical validation. *IEEE/ACM Transactions on Networking*, 8(2):133–145, April 2000.

P. Robert. On the asymptotic behavior of some algorithms. *Random Structures Algorithms*, to appear.

P. Tinnakornsrisuphap and R. La. Asymptotic behavior of heterogeneous TCP flows in RED gateway. draft, 2004.

Y. Yang and S. Lam. General AIMD congestion control. In *Proc. IEEE ICNP*, Osaka, Japan, November 2000.