Feedback and Bandwidth Sharing in Networks

Ayalvadi Ganesh, Peter Key and Laurent Massoulié
Microsoft Research, 7 J J Thomson Avenue, Cambridge CB3 0FB, UK
ajg, peterkey, lmassoul@microsoft.com

Abstract

Motivated by questions of bandwidth sharing in the Internet, we examine the problem of tracking a fair bandwidth allocation for a single link in the face of fluctuating user demands. We first use rate distortion theory to derive lower bounds on the variance of the tracking error as a function of the rate at which congestion feedback is relayed to users. We then combine the estimation and control problems. We derive the impact of feedback accuracy on performance when using either a Kalman filtering scheme or a simpler scheme proposed in [2].

1 Introduction and Preliminaries

The problem of sharing bandwidth between elastic users in a telecommunication network requires feedback from the network to the users about some measure of instantaneous congestion. For example, in the Internet, TCP uses packet drops caused by buffer overflow as the feedback mechanism. To what extent can the provision of more detailed feedback information by the network improve the rate allocation achieved by end users? We attempt to address this question by relating the accuracy with which user shares track a desired target to the rate at which they receive feedback.

We consider a slotted time model of N users sharing a single link. Users have different, time-varying utilities for bandwidth. The utility derived by user i from receiving bandwidth x in time slot t is given by $u_i(x) = w_i(t)\log x$. In keeping with terminology introduced by Gibbens and Kelly [2], we refer to $w_i(t)$ as user i's *willingness to pay* at time t[1]. If the aggregate load (bandwidth utilisation) on the link is x, then a cost $C(x)$ is incurred. The objective is to maximize the social welfare in each time slot:

$$V(\mathbf{x}(t)) = \sum_{i=1}^{N} u_i(x_i(t)) - C(x(t)), \tag{1}$$

where $\mathbf{x}(t)$ denotes the vector of bandwidth allocations to the individual users, and $x(t) = \sum_i x_i(t)$ denotes the aggregate bandwidth utilisation at time t. We denote by

[1] The choice of logarithmic utility functions is for ease of exposition. The results below can be derived in much greater generality; for example, the analysis extends easily to utility functions of the form $u_i(x) = w_i(t)^\alpha x^{1-\alpha}/(1-\alpha)$ for $\alpha > 0$. Logarithmic utility functions belong to this class and correspond to the limit $\alpha \to 1$.

$p(x)$ the marginal cost (or price) $C'(x)$of bandwidth at utilisation x. We assume that C is strictly increasing and convex or, equivalently, that p is positive and non-decreasing.

We can easily compute the optimal bandwidth allocations in (1); denoting these by $\mathbf{x}^*(t)$, we have

$$x_i^*(t) = \frac{w_i(t)}{p(x^*(t))} \quad \text{where} \;\; x^*(t)p(x^*(t)) = w(t). \tag{2}$$

Here, $x^*(t) = \sum_{i=1}^N x_i^*(t)$ and $w(t) = \sum_{i=1}^N w_i(t)$. A solution of the above equation exists if we assume that p is unbounded, and it is unique since p is non-decreasing[2].

How close can we come to the optimal solution above using a decentralized algorithm? In order to have a suitable metric of closeness, we consider the welfare loss due to deviation from optimality. Let $\mathbf{x}(t)$ be any rate allocation. If $\mathbf{x}(t)$ is close to $\mathbf{x}^*(t)$, we can use a Taylor series approximation to the welfare function in a neighbourhood of $\mathbf{x}^*(t)$. Recalling that $\nabla V(\mathbf{x}^*(t)) = 0$, we find that the loss of social welfare is given by

$$\begin{aligned} V(\mathbf{x}^*(t)) - V(\mathbf{x}(t)) &\approx \frac{-1}{2}\sum_{i=1}^N u_i''(x_i^*(t))(x_i(t)-x_i^*(t))^2 + \frac{1}{2}C''(x^*(t))(x(t)-x^*(t))^2 \\ &= \sum_{i=1}^N \frac{w_i(t)}{2}\left(\frac{x_i(t)-x_i^*(t)}{x_i^*(t)}\right)^2 + \frac{p'(x^*(t))}{2}(x(t)-x^*(t))^2. \end{aligned} \tag{3}$$

We seek to minimize the above quantity subject to constraints imposed by decentralisation and imperfect information.

First, we specify a stochastic model of the willingness-to-pay process. We model the evolution of the willingness-to-pay of each user as an autoregressive process:

$$w_i(t) = \rho w_i(t-1) + (1-\rho)\overline{w}_i + \overline{w}_i e_i(t), \tag{4}$$

where ρ is assumed to be the same for all users and $e_i(t)$ is a white Gaussian noise process, independent for different users. We denote by $w(t)$ the sum of $w_i(t)$ over all users. Hence,

$$w(t) = \rho w(t-1) + (1-\rho)\overline{w} + \overline{we}(t),$$

where $\overline{w}$ and $\overline{we}(t)$ denote the sum over i of $\overline{w}_i$ and $\overline{w}_i e_i(t)$ respectively.

This model can be motivated by thinking of the evolution of aggregate willingness to pay as arising from users entering or leaving the system (or going on or off). Then, the AR(1) model can be interpreted as an approximation of an $M/M/\infty$ model of the number of active users. Indeed, in the $M/M/\infty$ model where customers arrive at the points of a Poisson process of rate λ and depart after independent exponential holding times with mean $1/\mu$, the covariance between the number of customers present at time s and at time $t > s$ is given by $(\lambda/\mu)e^{-\mu(t-s)}$. If λ/μ, the mean number of customers in the system, is large, then the number of customers is well approximated by an Ornstein-Uhlenbeck process. An AR(1) model is obtained from this by periodic sampling. If we assume periodic sampling every δ time units, then by setting the AR(1) model parameters ρ and

[2] While we do not pursue game-theoretic aspects of bandwidth allocation in this paper, we note that the optimal bandwidth allocations above are close to being a Nash equilibrium, when the number of users N is large and no single user has a disproportionate impact on the system.

Var(e) to be $\rho = e^{-\mu\delta}$, and Var$(e) = (1-\rho^2)\lambda/\mu$, the covariance of the willingness to pay process matches that of the number of customers in the $M/M/\infty$ model.

In the next section, we focus on the communication issue and ignore the control aspect. We apply rate distortion and quantization theory to assess how accurately users can estimate the current value of an Ornstein-Uhlenbeck process given delayed feedback about past values.

In Section 3, we address control and communication jointly. We assess the impact on welfare loss of feedback accuracy when users employ Kalman filtering. In Section 4 we do the same for users employing a sub-optimal scheme proposed in [2]. We compare the structure of the simpler scheme to that of the optimal filter of Section 3.

2 Reconstruction error in the absence of control

2.1 Dedicated feedback

Let us first consider the situation where the resource (link) synthesises feedback individually for each user. Consider a specific user who receives feedback Y_n based on the price history before time $n\delta$, where δ is the frequency of feedback signals generation. Its task is then to reconstruct the price $p(t)$ at a given time t given the feedback received up to that time. Let us first consider reconstruction of $p(n\delta)$ given $\{Y_m, m \le n\}$. Consistent with (3) we use mean square error to assess the quality of the reconstruction: the distortion between the reconstructed signal $\hat{p}(n\delta)$ and $p(n\delta)$ is simply $\mathbf{E}(p(n\delta) - \hat{p}(n\delta))^2$. We make the assumption that the price process is of Ornstein-Uhlenbeck type, i.e.

$$dp(t) = \alpha(\bar{p} - p(t))dt + \sigma dB_t.$$

The sampled sequence $X_n := p(n\delta)$ is then a Gauss-Markov process, or an AR(1) process, described as

$$X_{n+1} = \rho X_n + (1-\rho)\bar{p} + Z_n,$$

where $\rho = e^{-\alpha\delta}$, $Z_n \sim \mathcal{N}(0, s)$, and $s = \sigma^2(1 - e^{-2\alpha\delta})/2\alpha$. We appeal to results of rate distortion theory for *sequential* reconstruction to assess the relation between optimal reconstruction and number of encoding bits per feedback symbol Y_n. Define the *sequential* rate distortion function $R(D)$ as

$$R(D) = \liminf_{n\to\infty} \frac{1}{n} \inf I(X_1^n; Y_1^n)$$

where the infimum is taken over all conditional distributions $p(Y_1^n|X_1^n)$ such that Y_1^m is independent of X_{m+1}^n conditionally on X_1^m, and $\mathbf{E}(X_m - Y_m)^2 \le D$, for all $m = 1, \ldots, n$. See [5] for background on this notion. It is shown in [5] that the rate $R(D)$ is a lower bound to the number of bits per symbol Y_n needed to reconstruct the original sequence X_n in a sequential manner, and with a distortion rate of D.

A computation of $R(D)$ for Gauss-Markov processes is performed in [5], yielding the formula

$$R(D) = \max\left(0, \frac{1}{2}\log\left(\rho^2 + \frac{s}{D}\right)\right),$$

or equivalently, the distortion rate function $D(R)$ is

$$D(R) = \frac{s}{2^{2R} - \rho^2}.$$

At a given time point t, the user will have access to the feedback signals Y_m such that $m\delta \leq t - \tau$, where τ is the round-trip propagation delay for that user. Assuming that $t - \tau$ is a multiple of δ, the corresponding distortion would then be,

$$D(R,\tau) = e^{-2\alpha\tau} D(R) + \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\tau}),$$

or equivalently

$$D(R,\tau) = \frac{\sigma^2}{2\alpha}\left(\frac{e^{-2\alpha\tau}(1 - e^{-2\alpha\delta})}{2^{2R} - e^{-2\alpha\delta}} + 1 - e^{-2\alpha\tau}\right).$$

The ratio of $D(\infty,\tau)$ to $D(0,\tau)$ is $1 - e^{-2\alpha\tau}$, so that feedback can prove useful only when $\alpha\tau$ is not too large, i.e. when the round-trip time τ is small compared to the critical time scale α^{-1} for the fluctuations in the price process.

Increasing the rate R will produce a measurable improvement when the fraction in the brackets is of the same order as the other terms in the bracket. This provides a rule of thumb for inferring when additional feedback will prove useful: increasing R will be worthwhile in the range $[0, R_c]$, where R_c is such that the two components in the bracket are equal, i.e.

$$R_c = \frac{1}{2}\log\left(e^{-2\alpha\delta} + \frac{e^{-2\alpha\tau}(1 - e^{-2\alpha\delta})}{1 - e^{-2\alpha\tau}}\right).$$

It can be seen that, when $\delta \leq \tau$, which we expect to be the normal situation, when more than one packet is sent in each round-trip time, the argument of the logarithm is maximized when either $\delta = \tau$ or $\delta = 0$, which yields

$$\delta < \tau \Rightarrow R_c \leq \frac{1}{2}\log(\max(1,2)) = \frac{1}{2}.$$

On the other hand, R_c can take on large values when $\tau < \delta$. In the case when $\alpha\tau < \alpha\delta << 1$,

$$R_c \approx \frac{1}{2}\log\left(1 + \frac{\delta}{\tau}\right).$$

These observations can be summarised as follows. Feedback appears to be useful only for connections with a round trip time τ that is of the same order as, or smaller than the critical time scale α^{-1} of price fluctuations. For feedback synthesised individually for each traffic source, having more than one bit of feedback per packet does not appear useful when more than one packet is sent per round trip time.

2.2 Anonymous feedback

We now assume that the resource is no longer aware of the identity of the traffic sources that generate each individual packet. The feedback signal attached to a given packet is thus generated based on the history of the global resource usage, but cannot rely on the history of feedback signals already sent to the corresponding traffic source. This is what we mean by anonymous feedback. We shall assume that the feedback corresponding to a given packet is a quantized version of the exact price at the time this packet is treated, with 2^R quantization levels. Quantization theory is surveyed in [3]. It is known for instance that a Gaussian random variable with variance s optimally quantized over 2^R levels is reconstructed from its quantized version with a mean square error approximately equal to $(\pi\sqrt{3}/2)s2^{-2R}$ for large R.

Suppose users do not exploit the full history of feedback they received, but only use the most recent signal. This removes the dependence of reconstruction error on sampling rate. For the same price process statistics as before, denoting by $D_q(R,\tau)$ the mean square distortion of the price predicted τ time units in advance by a traffic source, we thus have the approximate formula

$$D_q(R,\tau) \approx \frac{\sigma^2}{2\alpha}\left[\frac{\pi\sqrt{3}}{2}e^{-2\alpha\tau}2^{-2R} + 1 - e^{-2\alpha\tau}\right].$$

In the absence of feedback (i.e. for $R = 0$) the best estimate of price is its mean value, yielding $D_q(0,\tau) = \sigma^2/2\alpha$, and the ratio of $D_q(\infty,\tau)$ to $D_q(0,\tau)$ is again $1 - e^{-2\alpha\tau}$. Feedback thus proves useful only when $\alpha\tau$ is small.

The corresponding definition of R_c is

$$R_c = \frac{1}{2}\log\left(\frac{(\pi\sqrt{3}/2)e^{-2\alpha\tau}}{1 - e^{-2\alpha\tau}}\right).$$

Assuming that round trip times τ are of the order of 0.1 second, while α^{-1} is of the order of 100 seconds (think of the duration of a phone call) yields a value of $R_c = 5.2$.

These evaluations rely on optimal coding/decoding, which may not be practical in the context of the Internet where there are multiple bottlenecks. Indeed, the optimal quantization levels for a packet at a bottleneck would depend on the path of this packet through the network prior to reaching this link.

3 A Kalman filtering approach

We now consider the scenario where, in each time slot, the link computes its price based on the aggregate utilisation in that time slot. This information is quantized and fed back to users, possibly with some delay. Users have to choose their actions, i.e., their transmission rates, on the basis of the history of price feedback they have received.

In order to make the analysis tractable, we consider a linearisation of the willingness to pay, price and transmission rate processes around their mean values. Recalling the notation introduced in Section 1, let $\overline{x}$ solve $\overline{x}p(\overline{x}) = \overline{w}$, and let $\overline{x}_i = \overline{w}_i/p(\overline{x})$. We shall henceforth denote $p(\bar{x})$ by $\bar{p}$ and $p'(\bar{x})$ by $\bar{p}'$.

Define

$$v_i(t) = w_i(t) - \overline{w}_i,\ y_i(t) = x_i(t) - \overline{x}_i,\ y_i^*(t) = x_i^*(t) - \overline{x}_i,\ v(t) = w(t) - \overline{w},\ y(t) = x(t) - \overline{x}.$$

We can rewrite (4) as

$$v_i(t) = \rho v_i(t-1) + \overline{w}_i e_i(t). \tag{5}$$

Moreover, $(\overline{x} + y^*(t))p(\overline{x} + y^*(t)) = \overline{w} + v(t)$, which yields on linearisation that

$$[\bar{p} + \overline{x}\bar{p}']y^*(t) = v(t). \tag{6}$$

Using linearisation together with (6), we now obtain

$$y_i^*(t) = \frac{v_i(t)}{\bar{p}} - \frac{\overline{x}_i\bar{p}'}{\bar{p}}\frac{v(t)}{\bar{p} + \overline{x}\bar{p}'}. \tag{7}$$

Observe that $x_i(t) - x_i^*(t) = y_i(t) - y_i^*(t)$. Hence, we obtain from (3), (6) and (7), after neglecting higher-order terms, that

$$V(\mathbf{x}^*(t)) - V(\mathbf{x}(t)) = \sum_{i=1}^{N} \frac{\overline{w}_i}{2}\left[\frac{y_i(t)}{\overline{x}_i} - \frac{v_i(t)}{\overline{w}_i} + \frac{\overline{x}\bar{p}'}{\bar{p} + \overline{x}\bar{p}'}\frac{v(t)}{\overline{w}}\right]^2 + \frac{\overline{w}}{2}\frac{\overline{x}\bar{p}'}{\bar{p}}\left[\frac{y(t)}{\overline{x}} - \frac{\bar{p}}{\bar{p} + \overline{x}\bar{p}'}\frac{v(t)}{\overline{w}}\right]^2. \tag{8}$$

Let $\mathcal{F}_t^i$ denote the σ-algebra generated by the feedback information available to user i before time t. Clearly, the first term above is minimized by taking

$$y_i(t) = \frac{\overline{x}_i}{\overline{w}_i}v_i(t) - \frac{\overline{x}\bar{p}'}{\bar{p} + \overline{x}\bar{p}'}\frac{\overline{x}_i}{\overline{w}}\mathbf{E}[v(t)|\mathcal{F}_t^i].$$

If all users receive the same feedback information, so that $\mathcal{F}_t^i = \mathcal{F}_t$ does not depend on i, then the above choice of $y_i(t)$ yields

$$y(t) = \frac{1}{\bar{p}}\sum_{i=1}^{N} v_i(t) - \frac{\overline{x}\bar{p}'}{\bar{p} + \overline{x}\bar{p}'}\frac{\overline{x}}{\overline{w}}\mathbf{E}[v(t)|\mathcal{F}_t],$$

so that

$$\mathbf{E}[y(t)|\mathcal{F}_t] = \frac{\overline{x}}{\overline{w}}\left(1 - \frac{\overline{x}\bar{p}'}{\bar{p} + \overline{x}\bar{p}'}\right)\mathbf{E}[v(t)|\mathcal{F}_t],$$

and we see that the second term in (8) is minimized as well.

Define the quantity β via $\beta \stackrel{\text{def}}{=} \overline{x}\bar{p}'/[\bar{p} + \overline{x}\bar{p}']$ and note that $0 \le \beta \le 1$. We can now write the optimal filter equations as

$$y_i(t) = \frac{\overline{x}_i}{\overline{w}_i}\left\{v_i(t) - \beta\frac{\overline{w}_i}{\overline{w}}\mathbf{E}[v(t)|\mathcal{F}_t]\right\} \tag{9}$$

and

$$y(t) = \frac{1}{\bar{p}}\left\{v(t) - \beta\mathbf{E}[v(t)|\mathcal{F}_t]\right\} \tag{10}$$

If we choose $y_i(t)$ as above, then we obtain from (8) that the loss in social welfare is given by

$$V(\mathbf{x}^*(t)) - V(\mathbf{x}(t)) = \frac{\beta^2}{2(1-\beta)}\frac{1}{\overline{w}}\mathrm{Var}(v(t)|\mathcal{F}_t), \tag{11}$$

where, with an abuse of notation, $\mathrm{Var}(v(t)|\mathcal{F}_t) \stackrel{\text{def}}{=} \mathrm{Var}(v(t) - \mathbf{E}(v(t)|\mathcal{F}_{t-1}))$.

We suppose that the only information available to users at the beginning of time slot t is the history of prices up to the previous time slot, observed with quantization noise.

We take $v(t)$ to be the state of the system at time t. Its evolution is described by (4). We denote the observation at time t by $z(t)$ defined as

$$z(t) = \bar{p} + \bar{p}' y(t) + \eta(t) = \frac{\bar{p}'}{\bar{p}} v(t) + u(t) + \eta(t), \tag{12}$$

where $u(t) = \bar{p} - (\beta/\bar{p})\mathbf{E}[v(t)|\mathcal{F}_t]$ is an $\mathcal{F}_t$-measurable control, and $\mathcal{F}_t = \sigma(z(s), s \leq t-1)$. The noise term $\eta(t)$ models quantization error in encoding the price in a finite number of bits. We shall assume that $\eta(t)$ is an iid sequence of Gaussian random variables with mean zero and variance σ_η^2.

We are thus in the standard Kalman filtering setting. Let $P_t^- \stackrel{\text{def}}{=} \mathrm{Var}(v(t) - \mathbf{E}(v(t)|\mathcal{F}_t))$ denote the error variance of the state predictor at time t. This evolves according to a Riccati difference equation:

$$P_{t+1}^- = \overline{w}^2 {\sigma_e}^2 + \frac{\bar{p}^2 \rho^2 {\sigma_\eta}^2 P_t^-}{\bar{p}^2 {\sigma_\eta}^2 + \bar{p}'^2 P_t^-} \tag{13}$$

and converges to the unique limit P^- given by

$$\frac{P^-}{\overline{w}^2} = \frac{\sigma_e^2}{2} + \frac{-\left(1-\rho^2\right)\sigma_\eta^2 + \sqrt{4\,\gamma^2\,\bar{p}^2\,\sigma_e^2\,\sigma_\eta^2 + \left(\left(1-\rho^2\right)\,\sigma_\eta^2 - \gamma^2\,\bar{p}^2\,\sigma_e^2\right)^2}}{2\,\gamma^2\,\bar{p}^2}, \tag{14}$$

where $\beta = \overline{x}\bar{p}'/[\bar{p} + \overline{x}\bar{p}']$ and $\gamma = \beta/(1-\beta)$. We also have from (11) that

$$V(\mathbf{x}^*(t)) - V(\mathbf{x}(t)) = \frac{\beta^2}{2(1-\beta)} \frac{P^-}{\overline{w}}. \tag{15}$$

We now combine this with an expression for σ_η^2 suggested by the discussion of Section 2. We rely on the asymptotic quantization error formula for high rates and Gaussian quantized variables. In the present case, this reads

$$\sigma_\eta^2 = \frac{\pi\sqrt{3}}{2}\, 2^{-2R} \mathrm{Var}(p_t). \tag{16}$$

Applying linearisation, one has $p_t \approx \bar{p} + \bar{p}' y_t$. Also, standard arguments yield that

$$\mathrm{Var}(y_t) = \frac{1}{\bar{p}^2}\left[(1-\beta)^2 \frac{\sigma_e^2}{1-\rho^2} + P^-\right].$$

This equation combines with (14) and (16) to form a fixed point equation. This can be solved exactly as it is quadratic in the unknown σ_η^2. The curves in Figure 1 illustrate the dependency of the resulting quantity P^- as a function of ρ and R, for $\bar{x} = \bar{w} = \bar{p}' = 1$.
Special cases: If $\sigma_\eta^2 = 0$ or, equivalently, $R = \infty$, then the aggregate sending rate $x(t)$ is observed without error and the optimal predictive variance is

$$P^- = \overline{w}^2 \sigma_e^2.$$

At the other extreme, if $\sigma_\eta^2 = \infty$ there is no feedback information ($R = 0$), and the error is

$$\frac{\overline{w}^2 \sigma_e^2}{1-\rho^2}.$$

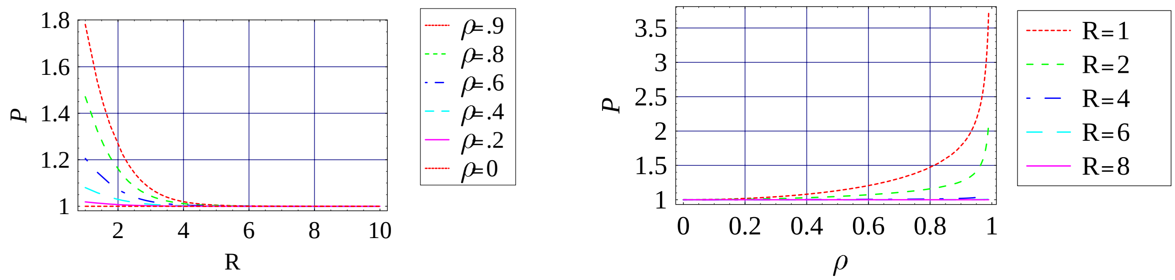

Figure 1: Normalized error as a function of number of bits R and ρ

The ratio of the two quantities is $1-\rho^2$ and is a measure of the value of feedback. This shows feedback has value only when ρ is close to 1, as expected in view of the discussion in Section 2.

Next, we describe the Kalman updating equation. The Kalman gain K is given by

$$K = \frac{\rho \bar{p}\, \bar{p}'\, P^-}{\bar{p}^2\, \sigma_\eta^2 + \bar{p}'^2 P^-}. \tag{17}$$

Writing $\hat{v}(t)$ as shorthand for $\mathbf{E}[v(t)|\mathcal{F}_t]$, we now have the update equation

$$\hat{v}(t+1) \;=\; \rho\hat{v}(t) + K\left(z(t) - \bar{p} - \frac{\bar{p}'}{\bar{p}}(1-\beta)\hat{v}(t)\right) \;=\; \rho\hat{v}(t) + K\frac{\bar{p}'}{\bar{p}}[v(t) - \hat{v}(t)] + K\eta(t).$$

Combining this with the description of x_i from equation (9) gives the recursion

$$\begin{aligned} x_i(t+1) - \rho x_i(t) \;&=\; \frac{w_i(t+1) - \rho w_i(t)}{\bar{p}} - \frac{\beta \overline{x}_i}{\overline{w}}[\hat{v}(t+1) - \hat{v}(t)] \\ &=\; \frac{w_i(t+1) - \rho w_i(t)}{\bar{p}} - K\frac{\beta \overline{x}_i}{\overline{w}}\left[\frac{\bar{p}'}{\bar{p}}(v(t) - \hat{v}(t)) + \eta(t)\right]. \end{aligned}$$

The term $w_i(t+1)/\overline{p}$ is the sending rate if the price is fixed at its mean value, so the recursion adds an adjustment that is a weighted sum of the previous difference between the rate sent and this quantity (at time t) and the difference between the observed price and its mean. The weights depend upon the rate at which information decays (ρ) and also on the gain parameters K and β, which each lie between 0 and 1.

4 The willingness-to-pay strategy

A simple scheme has been proposed by Gibbens and Kelly [2] and Kelly et al. [4] for rate adaptation by users that maximizes social welfare in equilibrium. This scheme is based on users comparing their willingness to pay with the charge incurred in each time slot, and adjusting their rates accordingly. The rate evolution is described by the equation

$$x_i(t+1) = x_i(t) + \kappa(w_i(t+1) - x_i(t)(p(t) + \eta(t))), \tag{18}$$

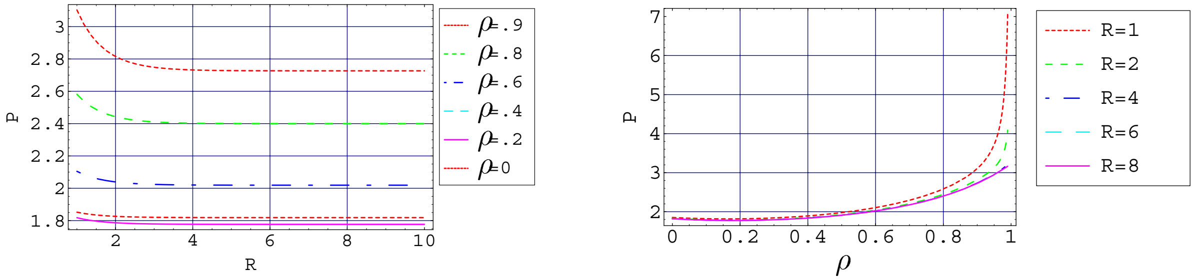

Figure 2: Welfare loss as a function of number of feedback bits R and correlation coefficient ρ, for $\kappa = 0.3$, with $w_i/\bar{w} = 0.01$, $\bar{p} = \bar{w} = \bar{p}' = 1$.

where x_i denotes the transmission rate of user i, w_i their willingness-to-pay, and p the price, which is a function of aggregate load. κ determines the speed of adjustment, and η is the noise with which the price information is fed back to users.

To compare the performance of the willingness to pay scheme with that of the framework described above, we now linearize around the stable point $\overline{x}_i\overline{p} = \overline{w}_i$. Defining the deviations from the stable point to be $y_i(t) = x_i(t) - \bar{x}_i$ and $y(t) = x(t) - \bar{x}$, gives the relation

$$y_i(t+1) = y_i(t) + \kappa(v_i(t+1) - y_i(t)\overline{p} - \overline{x}_i\overline{p}'y(t) - \overline{x}_i\eta(t)),$$

where $\overline{p}'$ is the derivative of the price function at the equilibrium rate $\overline{x}$. The variable $\xi(t) := (y_i(t), y(t), v_i(t+1), v(t+1))^T$ forms a vector AR(1) process,

$$\xi(t+1) = M\xi(t) + \zeta(t+1), \tag{19}$$

where

$$M = \begin{pmatrix} 1-\kappa\overline{p} & -\kappa\overline{x}_i\overline{p}' & \kappa & 0 \\ 0 & 1-\kappa(\overline{p}+\overline{x}\overline{p}') & 0 & \kappa \\ 0 & 0 & 1-\rho & 0 \\ 0 & 0 & 0 & 1-\rho \end{pmatrix},$$

and $\zeta(t+1) = (-\kappa\overline{x}_i\eta(t), -\kappa\overline{x}\eta(t), \bar{w}_i e_i(t+1), \bar{w}e(t+1))^T$. Let Σ denote the covariance matrix of $\zeta(t)$. We have

$$\Sigma = \begin{pmatrix} \kappa^2\overline{x}_i^2\sigma_\eta^2 & \kappa^2\overline{x}_i\overline{x}\sigma_\eta^2 & 0 & 0 \\ \kappa^2\overline{x}_i\overline{x}\sigma_\eta^2 & \kappa^2\overline{x}^2\sigma_\eta^2 & 0 & 0 \\ 0 & 0 & \bar{w}_i^2\sigma_e^2 & \bar{w}_i^2\sigma_e^2 \\ 0 & 0 & \bar{w}_i^2\sigma_e^2 & \bar{w}^2\sigma_e^2 \end{pmatrix}.$$

Straightforward but lengthy calculations enable us to find variances and covariances for $\xi(t)$. As in the case of Kalman filtering, we set the variance σ_η^2 as in (16), which again provides us with fixed point equations for σ_η^2. The dependence of the resulting loss in welfare on the number of feedback bits R and the correlation coefficient ρ are illustrated in Figure 2, with the parameter choices $\bar{w} = \bar{x} = \bar{p} = \bar{p}' = 1$, and $\bar{w}_i \equiv .01$ for all i.

5 Concluding remarks

We considered the problem of achieving weighted proportionally fair bandwidth allocations in a decentralized system with feedback. We derived information-theoretic lower bounds on the achievable mean square error of bandwidth allocations from their target values in terms of the feedback rate. For a model with time-varying weights and noisy feedback, we derived explicit expressions for the mean square error obtained by an optimal filtering scheme, and by a simpler scheme proposed in [2]. These results relate the mean square error to the rate at which the weights change and the variance of the noise in feedback. This variance can be related to the number of bits used to encode feedback information. The results show that the more slowly the weights (and hence the target bandwidth allocations) change over time, the greater the value of using multiple bits to provide more precise feedback.

References

[1] O.L.R. Jacobs, *Introduction to Control Theory*, Oxford University Press, 1993.

[2] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. Automatica (35) pp. 1969-1985, 1999.

[3] R.M. Gray and D.L. Neuhoff, Quantization, IEEE Trans. Info. Theory, 1998.

[4] F. Kelly, A. Maulloo and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. Journal of the Operational Research Society 49 pp. 237-252, 1998.

[5] S. Tatikonda, *Control under communication constraints*, PhD Thesis, MIT, 2001.