

Farsighted Users Harness Network Time-Diversity

Peter Key, Laurent Massoulié, Milan Vojnović

Microsoft Research Limited

7 JJ Thomson Avenue, Cambridge, UK

Email: {peterkey,lmassoul,milanv}@microsoft.com

Abstract—Fluctuations in network conditions are a common phenomenon. They arise in the current wired Internet due to changes in demand, and in wireless networks due to changing interference patterns. However, current congestion control design typically does not account for this, and in this sense the majority of congestion controllers proposed so far can be deemed as “myopic”. The present work deals with the following question: how should network end-users exploit such temporal fluctuations?

We introduce a formal framework, in which time diversity is explicitly described by phases in network condition. We propose as bandwidth allocation criterion the solution to an optimization problem, which features both classical (myopic) users and so-called farsighted users. We identify the corresponding farsighted user strategy as that maximizing throughput subject to a social norm related to TCP-friendliness. We establish basic desirable properties of the resulting allocations. We propose adaptive decentralized algorithms for farsighted users to achieve their target allocation. The algorithms do not require either explicit knowledge of dynamics in network conditions, or special feedback from the network.

I. INTRODUCTION

CONGESTION control of contending network end-users is traditionally assumed to aim at fair resource allocations. While the classical definition is that of max-min fairness (see [1]), the de-facto notion of fairness in force in the current Internet is that of TCP-fairness, i.e. allocations resulting from the use of standard TCP protocol stacks. Kelly and co-workers [7], [9], [4] proposed to use microeconomics formalism as a conceptual framework to discuss objectives of congestion control. According to this framework, the objective of congestion control is to maximize system welfare, defined as the aggregation of user utilities minus network costs.

There are two interpretations of this microeconomics formalism. The literal (and ambitious) interpretation requires users to be charged real money for their network usage, and react accordingly based on their utilities. According to the second, cautious interpretation, the framework provides a convenient generalization of available fairness criteria, each corresponding to a specific choice of utility functions (see [7], [14] for max-min

fairness, and [12], [8], [19] for TCP-fairness), and is useful for discussing properties of candidate congestion control algorithms, but should not be interpreted literally.

A limitation of this microeconomics framework is that it assumes fixed, static network conditions such as user populations or link capacities. Consequently, within this framework congestion controllers are required to match their marginal utility to the marginal cost of capacity along their path, *at any instant*. They are in this sense *myopic*, as they do not try to exploit potential time diversity in the network condition.

However, congestion control algorithms operate at time scales at which user populations vary. Hourly, predictable variations, are present in Internet traffic (see e.g. [13] for backbone measurements); less predictable fluctuations on time scales of seconds or minutes are also present. Wireless environments display fluctuations not only in competing user populations, but also in available capacities due to changing levels of interference, at time scales of the order of minutes.

Our goal in this paper is to investigate how congestion controllers could (and should) exploit such time diversity in network conditions. To this end, we extend the microeconomics framework of [9] to account for time diversity. In this extended framework, we introduce so-called *farsighted* users, whose aim is to maximize some utility function of their *average* throughput, in contrast with myopic users, whose aim is to maximize some utility function of their instantaneous rate at all times. As for the original framework of [9], our extended framework can be interpreted literally, assuming monetary transfers between network and users. This is appropriate to identify desirable user strategies for networks applying per-byte charging. An alternative application of the framework is to help reason about the implications of social norms such as TCP-friendliness [3], [6]. We find that greedy users that aim at maximizing their throughput while complying to a social norm of TCP-friendliness, are in fact a special case of our farsighted users.

Besides relating farsighted user strategies to TCP-friendliness, we establish the following structural properties. (i) The introduction of farsighted users tends to smooth network conditions, so that myopic users

see less temporal fluctuations. (ii) Farsighted strategy is evolutionarily stable, i.e. it is always advantageous for users to switch from myopic to farsighted, in that this increases their average throughput. Note that for real-time applications, average throughput may not be the adequate performance measure, and myopic strategies may be better suited; however delay-insensitive applications would benefit from switching to farsighted. (iii) The gain of switching to farsighted decreases with the number of users having switched. This can be interpreted as follows: the benefit derived from exploiting temporal diversity is shared out among farsighted users, hence the more they are, the smaller the benefit.

In the case of a single link scenario, we investigate more specific issues. In particular, we establish when farsighted users behave as low-priority users, in that they do not interfere with myopic users. For recent attempts to provide such low-priority service using end-point control only and without network support, see e.g. Key, Massoulié, and Wang [10], Kuzmanović and Knightly [2] Venkataramani, Kokku, and Dahlin [17]. We also characterize the impact of farsighted users on myopic users when these arrive at instants of a Poisson process, and stay in the system for the duration of a file transfer.

We address the problem of designing congestion control algorithms for farsighted users. We propose two classes of distributed algorithms. The first one gives an exact solution under a time scale separation assumption. The second one gives an approximate solution, but is easier to implement. These algorithms require neither specific network support, nor explicit identification of fluctuations in network conditions.

The outline of the paper is as follows. In Section 2 we describe the extended framework to account for time-diversity, and introduce farsighted users. In Section 3 we establish basic structural properties of the rate allocations achieved within this framework. In Section 4 we focus on the single link case. In Section 5 we describe the distributed algorithms for farsighted users. Section 6 contains numerical experiments. We conclude in Section 7.

a non smaller long-run time-average allocation, than a competing persistent myopic user, with the same utility function (resp. larger, if the farsighted user has zero allocation for a phase).

II. MICROECONOMICS FRAMEWORK FOR RATE ALLOCATIONS UNDER TIME-DIVERSITY

Let us consider a network consisting of a set \mathcal{L} of links ℓ , and a set \mathcal{R} of users r . We assume that the network goes through phases, and denote by $u(t)$ the phase at time t . Phases are assumed to belong to a finite set \mathcal{U} . In addition, we assume that the process $u(t)$ is

SYSTEM:

$$\begin{aligned} & \text{maximize} && \sum_{r \in \mathcal{R}} \underline{U}_r(x_r) - \\ & && - \sum_{u \in \mathcal{U}} \pi(u) \sum_{\ell \in \mathcal{L}} C_{\ell,u} \left(\sum_q A_{q\ell}(u) x_q(u) \right) \\ & \text{over} && x_r(u) \geq 0, \quad r \in \mathcal{R}_u, \quad u \in \mathcal{U}. \end{aligned}$$

Fig. 1. Microeconomics framework for networks with time-diversity. Time-diversity is represented by distinct phases $u \in \mathcal{U}$ and the corresponding fraction of time $\pi(u)$.

given exogenously, and is ergodic. We denote by $\pi(u)$ the proportion of time spent in phase u .

The network condition phase determines two aspects: (i) the population of users competing for resources depends on the current phase; we note \mathcal{R}_u the population of users active in phase u . (ii) The cost incurred at link ℓ for carrying traffic at rate x , which we write $C_{\ell,u}(x)$, also depends on u . (iii) The route taken by traffic from user r is determined by the phase. The routing structure is summarized in the routing matrix $A_{r\ell}(u) = 1$ if data from user r is routed through link ℓ during phase u , and $A_{r\ell}(u) = 0$ otherwise.

In order to complete the specification of the framework, for each user r , we assume there is a utility function \underline{U}_r mapping $\mathbb{R}_+^{\mathcal{U}}$ to \mathbb{R}_+ , and such that the utility to user r of sending at rate $x_r(u)$ in each phase $u \in \mathcal{U}$ is given by $\underline{U}_r(x_r)$, where $x_r = \{x_r(u)\}_{u \in \mathcal{U}}$. Note that if $r \notin \mathcal{R}_u$ we necessarily set $x_r(u) = 0$. We finally assume that the utility functions \underline{U}_r are concave, and the cost functions $C_{\ell,u}$ are convex.

Under these assumptions, we consider as the system objective to maximize the total aggregated welfare, i.e. solve the SYSTEM optimization problem in Figure 1.

The above framework can in fact be formally interpreted as an instance of the network resource allocation framework with multi-path routing considered by Gibbens and Kelly [5]. Indeed in their context, users' utilities are functions of each rate that they can use along each route available to them. By viewing each pair (ℓ, u) as a virtual link with associated cost function $\pi(u)C_{\ell,u}$, and by interpreting $x_r(u)$ as the rate sent by user r along its route labeled u , which consists of the links (ℓ, u) such that $A_{r\ell}(u) = 1$, we indeed have a special instance of their framework. The main difference thus lies in the interpretation, as we have temporal rather than spatial diversity. This will have significant implications on what algorithms can be used to solve the problem.

We now specify two types of users that we shall be interested in the sequel. We shall consider, on the one hand, *farsighted* users, whose utility function \underline{U}_r reads

$$\underline{U}_r(x_r) = U_r(\bar{x}_r),$$

where U_r is some utility function mapping \mathbb{R}_+ to \mathbb{R}_+ , and \bar{x}_r is the average rate obtained by user r , given by

$$\bar{x}_r = \sum_{u \in \mathcal{U}} \pi(u) x_r(u), \quad r \in \mathcal{R}.$$

We denote by \mathcal{F} the set of all farsighted users.

On the other hand, we shall consider *myopic* users, whose utility function reads

$$\underline{U}_r(x_r) = \sum_{u \in \mathcal{U}} \pi(u) U_r(x_r).$$

Here as for farsighted users, U_r denotes some utility function mapping \mathbb{R}_+ to itself.

It should be intuitively clear why such users are called myopic: if the user population consists of myopic users only, then the global SYSTEM separates in independent SYSTEM problem instances, one for each phase u . Thus when solving SYSTEM, such users choose their sending rates in each phase independently of the network conditions in the other phases.

We use the following terminology in the sequel. We define the *marginal cost* at link ℓ in phase u to be the derivative of the cost function $C_{\ell,u}$ evaluated at its current load, and denote it by $p_\ell(u)$. According to the literal interpretation of the framework, this represents an actual price per byte; in the alternative interpretation, this is a non-monetary signal, such as a loss rate.

III. BASIC PROPERTIES

This section describes several properties of farsighted users, and of their interaction with myopic users.

A. Price equalization and smoothing

The SYSTEM allocation for farsighted users is characterized by the following property.

Theorem 1: For a farsighted user $r \in \mathcal{F}$, there exists a positive $p_r^* > 0$ such that $p_r^* = U_r'(\bar{x}_r)$, and for all $u \in \mathcal{U}$,

$$\begin{aligned} p_r(u) &= p_r^* && \text{if } x_r(u) > 0, \\ p_r(u) &\geq p_r^* && \text{otherwise,} \end{aligned}$$

where $p_r(u) := \sum_{\ell \in \mathcal{L}} A_{r\ell}(u) C'_{\ell,u}(\sum_q A_{q\ell}(u) x_q(u))$.

This tells us that the farsighted user r equalizes prices $p_r(u)$ across good network phases, $\mathcal{G}_r := \{u \in \mathcal{U} : x_r(u) > 0\}$, where “good” is as perceived by the user r . The farsighted user r has a positive allocation only for good phases \mathcal{G}_r , else, for bad phases, the user receives no allocation. In this sense, it smoothers network conditions: myopic users taking the same network route as user r will send at the same rate in all good phases, as they adapt their rate in phase u to the corresponding marginal cost $p_r(u) \equiv p_r^*$.

Proof: The property follows directly from the first-order optimality conditions of SYSTEM:

$$U_r'(\bar{x}_r) = p_r(u) - v_r(u), \quad r \in \mathcal{F}$$

where $v_r(u)$ is the Lagrange multiplier associated to the non-negativity constraint $x_r(u) \geq 0$. As such, it satisfies $v_r(u) \geq 0$ and $v_r(u) x_r(u) = 0$. The result follows. ■

B. Conservativeness and throughput maximization

The average marginal cost for a farsighted user r reads $\bar{p}_r = [\sum_{u,\ell} \pi(u) A_{\ell r}(u) p_{\ell u} x_r(u)] / [\sum_u \pi(u) x_r(u)]$; by the above price equalization result, this also reads $\bar{p}_r^* = U_r'(\bar{x}_r)$. Consider then the following *conservativeness* requirement imposed to user r : its average rate \bar{x}_r and its average marginal cost \bar{p}_r have to satisfy

$$\bar{p}_r \leq U_r'(\bar{x}_r), \quad (1)$$

where the utility function U_r is mandated by the system designer. This notion of conservativeness has been identified in [18] and proposed as an alternative to the notion of TCP friendliness.

Upon writing the first-order optimality conditions for SYSTEM, one sees that farsighted users achieve equality in (1). Because the function U_r' is decreasing, it follows that farsighted users aim at maximizing throughput, under the conservativeness relation (1).

C. Evolutionary stability of farsighted against myopic

We now show that farsighted users retrieve more from the system than their myopic counterparts.

Theorem 2: Consider a farsighted user f and a persistent myopic user m that compete for the same set of resources. Suppose both users have the same utility function, i.e. $U_f = U_m$. Then it holds that

$$\bar{x}_m \leq \bar{x}_f. \quad (2)$$

Proof: Let \mathcal{G}_f be the set of good phases for the user f . From the first-order optimality conditions for SYSTEM(U, C), we have

$$U_f'(\bar{x}_f) = U_m'(x_m(u)), \quad u \in \mathcal{G}_f \quad (3)$$

$$U_f'(\bar{x}_f) \leq U_m'(x_m(u)), \quad u \notin \mathcal{G}_f. \quad (4)$$

Because $U_m = U_f$, it follows that $x_m(u) = \bar{x}_f$ for all $u \in \mathcal{G}_f$. Thus,

$$\begin{aligned} \bar{x}_m &= \left(\sum_{u \in \mathcal{G}_f} \pi(u) \right) \bar{x}_f + \sum_{u \in \mathcal{U} - \mathcal{G}_f} \pi(u) x_m(u) \\ &= \bar{x}_f - \sum_{u \in \mathcal{U} - \mathcal{G}_f} \pi(u) (\bar{x}_f - x_m(u)). \end{aligned}$$

From (4), the last sum is non-negative which yields the asserted inequality. ■

This result entails that a user interested in average throughput maximization would benefit more from using the farsighted strategy than the myopic one, given that the utility functions U_f and U_m are the same. The result does not say the average throughput of a farsighted user is larger than that of a competing myopic user, but note that in situations when there exists a phase such that strict inequality obtains in (4), then strict inequality in (2) holds. Thus, given a choice, users would switch from myopic to farsighted.

D. Diminishing returns of switching strategies

We now consider the impact on farsighted users of myopic users switching to farsighted strategy. We shall see that the average throughput to farsighted users diminishes as new users adopt the farsighted strategy. The interpretation is that the benefits of using this strategy become shared out among more users, hence there is a diminishing return in switching as the farsighted population increases. Formally, we have

Theorem 3: Consider a system of n users utilizing the same network routes, and among which k are farsighted, and $n - k$ are myopic. Denote by $\bar{x}_f(k)$ the average throughput to farsighted users as a function of k . Then $\bar{x}_f(k)$ decreases with k .

Proof: The rates $x_m(u)$, $x_f(u)$ obtained by myopic and farsighted users respectively solve the problem:

$$\begin{aligned} & \text{maximize} && kU(\bar{x}_f) + (n-k) \sum_{u \in \mathcal{U}} \pi(u)U(x_m(u)) \\ & && - \sum_{u \in \mathcal{U}} \pi(u)C_u(kx_f(u) + (n-k)x_m(u)) \\ & \text{over} && x_r(u) \geq 0, \quad r \in \{f, m\}, \quad u \in \mathcal{U}, \end{aligned}$$

where C_u denotes the aggregated cost function along the route taken by users in phase u . First-order optimality conditions imply the following relation:

$$x_f(u) = \frac{1}{k} [C'_u{}^{-1}(U'(\bar{x}_f)) - (n-k)\bar{x}_f]^+. \quad (5)$$

Denote by g_u the function $C'_u{}^{-1}(U'(\cdot))$. Differentiating this identity with respect to k , we obtain

$$\frac{\partial x_f(u)}{\partial k} = -\frac{1}{k}x_f(u) + \frac{1}{k} \mathbb{1}_{\{x_f(u) > 0\}} \left[(g'_u(\bar{x}_f) - (n-k) \frac{\partial \bar{x}_f}{\partial k} + \bar{x}_f) \right].$$

Multiplying this identity by $\pi(u)$ and summing over u , after elementary manipulations we arrive at

$$(1+A) \frac{\partial \bar{x}_f}{\partial k} = -\frac{\bar{x}_f}{k} \sum_{u \in \mathcal{U}} \mathbb{1}_{\{x_f(u)=0\}} \pi(u), \quad (6)$$

where $A = \frac{1}{k} \sum_{u \in \mathcal{U}} \mathbb{1}_{\{x_f(u) > 0\}} [(n-k) - g'_u(\bar{x}_f)]$.

By concavity of U and convexity of C_u , g'_u is non-positive, hence $A > 0$. The previous identity then implies the announced result. ■

One may be tempted proposing that under the setting as in Theorem 3, the average throughput of a myopic user would increase as the fraction of farsighted users increases, but this is not necessarily true (proof in Appendix A):

Proposition 1: Under same setting as in Theorem 3, the average rate of a myopic user $\bar{x}_m(k)$ is not *always* increasing with k .

In summary, we showed: the larger the fraction of farsighted users in a fixed number of users sharing the same set of routes, the smaller the average throughput of a farsighted user. The average throughput of a myopic user may increase, but not always.

IV. FURTHER PROPERTIES FOR THE SINGLE LINK SCENARIO

In this section we confine our attention to the case of a single link. We first characterize optimum allocation, which we then use in obtaining an order relation for mean number of myopic users for a system with queueing, and finally relate farsighted and low-priority users.

A. Optimum allocation for single resource

Locally to this subsection consider a link with fixed capacity c . We assume no cost is incurred for utilizations below c . A population of non-persistent myopic users compete with a farsighted user. The time-diversity is due to fluctuations of the number of non-persistent myopic users that arrive and leave the resource. We assume the respective utility functions of the farsighted and a myopic user are U_f and U_M . A phase $u(t)$ is now interpreted as the number of myopic users in the system at time t .

Theorem 4: SYSTEM optimum allocation is

$$\begin{aligned} x_M(u) &= \begin{cases} cx^* & u \leq u^* \\ \frac{c}{u} & \text{else,} \end{cases} \\ x_f(u) &= \begin{cases} c(1-x^*u) & u \leq u^* \\ 0 & \text{else,} \end{cases} \end{aligned}$$

where u^* is largest $u \in \mathcal{U}$ such that $ux^* < 1$ and x^* is a solution of

$$U'_M(x) = U'_f(\sum_{u \in \mathcal{U}} \pi(u)(1-ux)^+). \quad (7)$$

Corollary 1: Under $U_f \equiv U_M$, (7) reads as

$$x = \sum_{u=0}^{u^*} \pi(u)(1-ux). \quad (8)$$

where

$$u^* = \max \{u \in \mathcal{U} : \sum_{n=0}^u (u-n)\pi(n) < 1\}.$$

Figure 2 illustrates solving the implicit function (8).

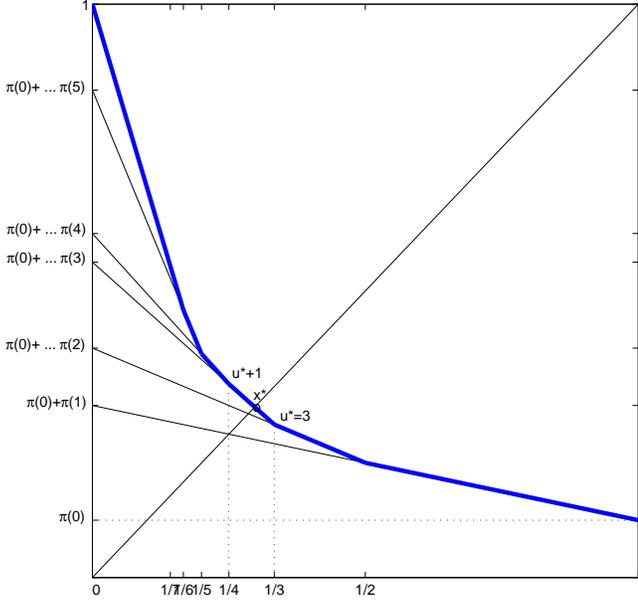


Fig. 2. The thick curve shows the function that acts as argument to $U'_{\mathcal{F}}(\cdot)$ in equation (7). On an interval $[1/(u+1), 1/u]$, the function is $x \rightarrow \sum_{n=0}^u \pi(n) - (\sum_{n=1}^u \pi(n)n)x$. Under $U_{\mathcal{M}} \equiv U_{\mathcal{F}}$, the intersection of this function and the line $x \rightarrow x$ is the solution of (7).

The interpretation is a specific instance of the price equalization property established Section III-A. The farsighted user receives a positive rate at times when the number of myopic users is not larger than the threshold u^* (the threshold defines good phases = number of myopic users not larger than u^*). When in a good phase, a myopic user receives a fixed rate x^* , whereas the rate of the farsighted user decreases linearly with the number of competing myopic users in the good phases. In a bad phase (= number of competing myopic users larger than the threshold u^*), the farsighted user receives nil rate; the myopic users share the resource as they would in absence of the farsighted user.

B. Farsighted induces smaller sojourn time for myopic

It appears a plausible claim that the farsighted backing off when phases are bad may be beneficial with respect to the sojourn time in the system of competing myopic users. The claim is motivated by the intuition that as the number of competing myopic users accumulates, the system turns from good to bad phases, and then a competing farsighted user backs off and let the competing myopic users leave the system. In this section, the claim is proved under some special assumptions.

With no loss of generality consider a link with unit capacity. Suppose myopic users arrive according to homogeneous Poisson process with rate λ and each arrival is a file transfer with exponentially distributed file size with mean $1/\mu$. For stability, we require $\rho := \lambda/\mu <$

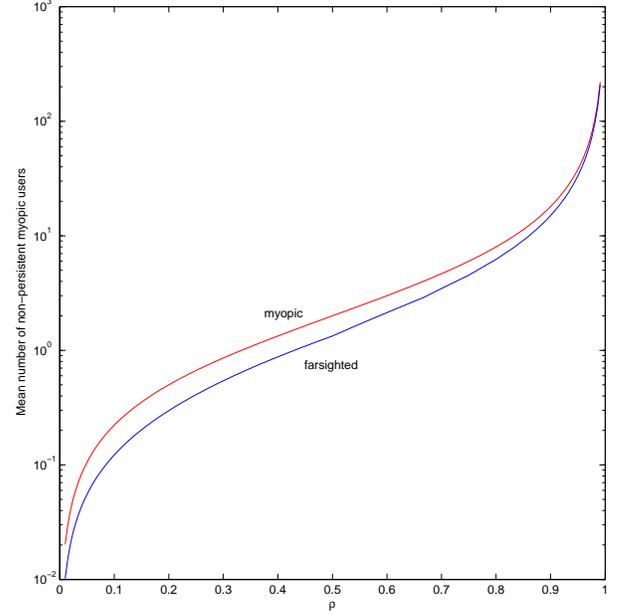


Fig. 3. The mean number of non-persistent myopic users competing with: (top curve) persistent myopic user, and (bottom curve) persistent farsighted user. The farsighted user induces smaller mean number of non-persistent myopic users in the system, than if the persistent user is myopic.

1, and to exclude a degenerate case of no interest, assume $\rho > 0$. We assume all users have identical utility functions. We have the following result:

Proposition 2: Let $\bar{n}^{\mathcal{F}}$ [resp. $\bar{n}^{\mathcal{M}}$] be the mean number of non-persistent myopic users when the persistent user is farsighted [resp. myopic]. Under the prevailing assumptions, the following relation is true: for any $0 < \rho < 1$,

$$\bar{n}^{\mathcal{F}} < \bar{n}^{\mathcal{M}}.$$

The relation entails (by Little's law) that non-persistent myopic users have smaller mean sojourn time, and hence larger throughput, if they compete with a persistent farsighted user, than if the persistent user is myopic. The ordering is complemented with numerical values in Figure 3.

In the remainder of this section we proceed by analysis of the two systems that yields $\bar{n}^{\mathcal{F}}$ and $\bar{n}^{\mathcal{M}}$, and then ultimately show that the asserted ordering is true.

Case 1: persistent user is farsighted. Consider the system when the persistent user is farsighted. The number of the myopic users present in the system evolves as a birth-and-death Markov process with transition rates

$$\begin{aligned} u \rightarrow u+1 & \text{ with rate } \lambda \\ u \rightarrow u-1 & \text{ with rate } \mu(ux^*1_{u \leq u^*} + 1_{u > u^*}). \end{aligned}$$

The stationary distribution π is

$$\pi(u) = \pi(0) \begin{cases} \frac{1}{u!} \left(\frac{\rho}{x^*}\right)^u & u \leq u^* \\ \frac{1}{x^{*u} u^*!} \rho^u & u > u^*, \end{cases}$$

where

$$\pi(0) = \frac{1}{\sum_{n=0}^{u^*-1} \frac{1}{n!} \left(\frac{\rho}{x^*}\right)^n + \frac{1}{1-\rho} \frac{1}{u^*!} \left(\frac{\rho}{x^*}\right)^{u^*}}. \quad (9)$$

From (8), an anticipated relation is easily obtained

$$x^* = 1 - \rho.$$

By definition of u^* , we have

$$\frac{1}{1-\rho} - 1 \leq u^* < \frac{1}{1-\rho}. \quad (10)$$

Hence

$$u^* = \left\lfloor \frac{1}{1-\rho} \right\rfloor. \quad (11)$$

It can be readily checked that the mean number of myopic users is

$$\bar{n}^{\mathcal{F}} = \frac{\rho}{1-\rho} \left[1 + \pi(0) \frac{1}{(u^*-1)!} \left(\frac{\rho}{1-\rho}\right)^{u^*} \right]. \quad (12)$$

Case 2: persistent user is myopic. Consider the system with myopic persistent user. The evolution of the number of non-persistent myopic users obeys the dynamics

$$\begin{aligned} u &\rightarrow u+1 && \text{with rate } \lambda \\ u &\rightarrow u-1 && \text{with rate } \mu \frac{u}{u+1}. \end{aligned}$$

The stationary distribution π and mean are readily obtained

$$\begin{aligned} \pi(u) &= (1-\rho)^2 (u+1) \rho^u, \\ \bar{n}^{\mathcal{M}} &= \frac{2\rho}{1-\rho}. \end{aligned} \quad (13)$$

Proof: [Proposition 2] From (12) and (13) note that the asserted inequality holds if and only if

$$\pi(0) \frac{1}{(u^*-1)!} \left(\frac{\rho}{1-\rho}\right)^{u^*} < 1.$$

By plugging (9), the last relation can be re-written as

$$\begin{aligned} &\frac{1}{(u^*-1)!} \left(1 - \frac{1}{u^*(1-\rho)}\right) \left(\frac{\rho}{1-\rho}\right)^{u^*} \\ &< \sum_{n=0}^{u^*-1} \frac{1}{n!} \left(\frac{\rho}{1-\rho}\right)^n. \end{aligned} \quad (14)$$

From the right inequality in (10), we have that $u^*(1-\rho) < 1$, hence, the left-hand side in (14) is smaller than 0. Now, note that the right-hand side in (14) is larger or equal to 1 ($u^* \geq 1$, which follows from the left inequality in (10) and the assumption that $\rho > 0$), so the relation in (14) always holds. ■

C. When farsighted users are low-priority background

It is of interest to know whether a population of farsighted users can act as low-priority background to myopic users competing for a set of resources. We say a farsighted user is low-priority background if the user receives a positive allocation only at times when no myopic user competes for a resource that the farsighted user intends to use. In this section we provide a necessary and sufficient condition for this to hold for arbitrary cost functions.

Theorem 5: Consider a single resource with an associated increasing, convex cost function C . A population of f farsighted users competes with non-persistent myopic users. The farsighted users have a concave utility function $U_{\mathcal{F}}$ and that of the myopic users is $U_{\mathcal{M}}$, assumed to be strictly concave. The farsighted population is low-priority background for any choice of an increasing convex function C , if and only if:

$$(LP): \quad U'_{\mathcal{F}} \left(\frac{\pi(0)}{f} x \right) \leq U'_{\mathcal{M}}(x), \quad \text{all } x > 0.$$

The proof is deferred to the appendix. This result suggests a farsighted user can be made low-priority background to myopic users if the load of myopic users is kept sufficiently low, and if suitable utility functions are chosen. For instance, suppose the utility functions $U_{\mathcal{F}}$ and $U_{\mathcal{M}}$ belong to the family of α -fair utility functions,

$$U_r(x) = w_r \frac{x^{1-\alpha}}{1-\alpha}, \quad \alpha > 0. \quad (15)$$

Now, (LP) reads as $\frac{w_{\mathcal{F}}}{w_{\mathcal{M}}} \leq \left(\frac{\pi(0)}{f}\right)^{\alpha}$.

V. ALGORITHMS

As we mentioned previously, our general framework is a special instance of that in [5]. It is thus tempting to try and re-use the algorithms described in [5] for solving their multi-path allocation problem to solve our target SYSTEM problem. However, the congestion controllers obtained in this way would require as an input an explicit description of phases $u \in \mathcal{U}$, knowledge of the proportion of time spent in each phase, and knowledge of the current phase at any given time. This is not practically feasible. We thus propose new algorithms for solving our SYSTEM problem, which require none of the above. The first class of algorithms solves the SYSTEM problem exactly, under a time scale separation assumption. The second provides only an approximate solution, but is better suited for practical implementation.

A. Exact algorithms under time scale separation

We consider three time scales of interest: the shortest corresponds to rate adaptation of users. It is dictated by

round-trip propagation delays, and thus, in the context of the Internet, is of the order of tens to hundreds of milliseconds. The second time scale corresponds to that at which network conditions change, i.e. this is the time scale of fluctuations in either user populations, route selections, and interference levels in wireless scenarios. We assume this is of the order of minutes. Finally, the longest time scale we consider is that over which farsighted users measure their own performance, i.e. this is the time over which the long-term average rates \bar{x}_r are computed. We may assume these are of the order of hours. We refer to the three time scales as τ_1 , τ_2 and τ_3 . We assume that $\tau_1 \ll \tau_2 \ll \tau_3$.

We shall use the following notations. \mathcal{F} (respectively, \mathcal{M}) denotes the set of farsighted (respectively, myopic) users. In phase $u \in \mathcal{U}$, the variable $n_r(u) \in \{0, 1\}$ represents whether user r is active or not. Note that for persistent users (and in particular, farsighted users), we have $n_r(u) \equiv 1$.

Consider the following algorithm for farsighted users. User r maintains an average measure of marginal utility, denoted by ξ_r . On short time scales, it aims to maximize its net benefit evaluated not on the basis of its actual utility function U_r , but rather based on the basis of the linear utility function $x_r \rightarrow \xi_r x_r$. Thus, in phase u , the instantaneous optimization problem users try to solve is the following:

$$\begin{aligned} \text{maximize} \quad & \sum_{r \in \mathcal{F}} x_r \xi_r + \sum_{q \in \mathcal{M}} n_q(u) U_q(x_q) - \\ & - \sum_{\ell \in \mathcal{L}} C_{\ell, u} \left(\sum_j A_{j\ell}(u) n_j(u) x_j \right). \end{aligned} \quad (16)$$

We assume that all users are able to select rates x_r that maximize the above instantaneous objective function, in a time scale of order τ_1 . This is feasible for instance provided all users receive instantaneous feedback on the aggregate marginal cost along their route through the network, conveyed for instance via ECN marks carried by ACK packets, or loss events, and if users adapt their sending rates based on such feedback according to schemes such as those in [9]. For definiteness, consider the following:

$$\boxed{\frac{dx_r}{dt} = \begin{cases} \kappa_r (\xi_r - \sum_{\ell} A_{r\ell}(u) q_{\ell}) & \text{if } r \in \mathcal{F}, \\ \kappa_r (U'_r(x_r) - \sum_{\ell} A_{r\ell}(u) q_{\ell}) & \text{if } r \in \mathcal{M}, \end{cases}} \quad (17)$$

where the gain parameters κ_r are of the order of $1/\tau_1$, i.e. the reciprocal of the shortest time scale, and

$$q_{\ell} = C'_{\ell} \left(\sum_j A_{j\ell} n_j(u) x_j \right). \quad (18)$$

It may be readily checked, using LaSalle's invariance principle (see e.g. [11]) that these dynamics do converge

to the set of maximizers of the objective function in (16), as this objective function is a Lyapunov function for the dynamics described here.

We thus assume that, seen from time scale τ_2 , user rates always lie in the set of maximizers of (16). We adopt a more compact notation by introducing the function

$$\Gamma_u((x_i)_{i \in \mathcal{F}}) := - \sup_{(x_j)_{j \in \mathcal{M}}} \left\{ \sum_{j \in \mathcal{M}} n_j(u) U_j(x_j) - \sum_{\ell \in \mathcal{L}} C_{\ell, u} \left(\sum_i A_{i\ell}(u) n_i(u) x_i \right) \right\}.$$

It is readily seen that this is a convex function. We then have the characterization of x_r , $r \in \mathcal{F}$ in terms of the vector ξ of average marginal utilities and the subgradient of Γ_u :

$$\xi \in \partial \Gamma_u(x). \quad (19)$$

We now use Theorem 23.1, p. 213-214 in [16] to transform this characterization into one that will prove useful in establishing optimality properties of our proposed adaptation scheme. Recall that the (Legendre-Fenchel) convex conjugate Γ_u^* of Γ_u is defined by

$$\Gamma_u^*(y) = \sup_x \{ \langle x, y \rangle - \Gamma_u(x) \}.$$

According to Theorem 23.1, p. 213-214 in [16], under standard technical assumptions¹, we have the equivalent characterization of (19):

$$x \in \partial \Gamma_u^*(\xi). \quad (20)$$

We now complete our specification of the control algorithms used by the farsighted users. They adapt their parameters ξ_r according to:

$$\boxed{\frac{d\xi_r}{dt} = \alpha_r \left(U'_r(x_r) - \xi_r \right).} \quad (21)$$

Here, the gain parameter α_r is of the order of $1/\tau_3$, i.e. the reciprocal of the largest time scale. The main appeal of this algorithm is that it does not require to explicitly monitor the fluctuations in the phases u ; these are accounted for implicitly.

Let us first see why the corresponding algorithm achieves the desired optimum, in the special case where there are no fluctuations in the phase u .

¹Namely, that the convex function Γ_u be proper and lower semi-continuous (see definitions in [16], p. 24 and 51 respectively). This holds if the functions U_j and $C_{\ell, u}$ have these properties.

Lemma 1: For fixed u , the dynamics specified by (19) (or equivalently, (20)) and (21) admit the Lyapunov function

$$L(\xi) = \sum_{r \in \mathcal{F}} \int_0^{\xi_r} U_r'^{-1}(v) dv - \Gamma_u^*(\xi), \quad (22)$$

in that the value of $L(\xi)$ increases over time. Moreover, $\xi(t)$ converges to the set of maximizers of L , and maxima in ξ of L are in one to one correspondence to maxima in x of the original objective function

$$\sum_{r \in \mathcal{R}} n_r(u) U_r(x_r) - \sum_{\ell \in \mathcal{L}} C_{\ell,u} (\sum_r A_{r\ell}(u) n_r(u) x_r)$$

via the relation (19).

Proof: In view of (21) and (20), $d\xi_r/dt$ has the same sign as

$$U_r'^{-1}(\xi_r) - x_r = U_r'^{-1}(\xi_r) - (\partial \Gamma_u^*(\xi))_r.$$

Thus, we see that the value of $L(\xi)$ increases strictly over time, except at those points where the right-hand side in the above equation is zero for all r . However, this can hold only at maxima of the concave objective function L . It thus only remains to establish the correspondence of such maxima in ξ to maxima in x of the original objective function via (19). This readily follows by combining the identities $\xi_r = U_r'(x_r)$ with (19) and then recalling the definition of Γ_u . ■

Let us now deal with time-varying phases $u(t)$. For convenience we recall the target SYSTEM optimization problem:

$$\text{maximize} \quad \sum_{r \in \mathcal{F}} U_r(\bar{x}_r) + \quad (23)$$

$$\begin{aligned} &+ \sum_{u \in \mathcal{U}} \pi(u) \left(\sum_{q \in \mathcal{M}} n_q(u) U_q(x_q(u)) - \right. \\ &\left. - \sum_{\ell \in \mathcal{L}} C_{\ell,u} \left(\sum_{i \in \mathcal{F} \cup \mathcal{M}} A_{i\ell} n_i(u) x_i(u) \right) \right) \end{aligned} \quad (24)$$

$$\text{subject to} \quad \bar{x}_r = \sum_{u \in \mathcal{U}} \pi(u) x_r(u), \quad r \in \mathcal{F}, \quad (25)$$

$$\text{over} \quad x_r(u) \geq 0, \quad r \in \mathcal{F} \cup \mathcal{M}, \quad u \in \mathcal{U}. \quad (26)$$

We now give a heuristic derivation of a characterization of the dynamics of ξ when time scale separation $\tau_2 \ll \tau_3$ holds. Namely, we argue that dynamics of ξ at time scale τ_3 are again specified by (21), but the characterization of x_r in (20), which specifies constraints on x_r for any value of u , should be replaced by the corresponding ‘‘averaged’’ characterization:

$$x(t) \in \partial \left\{ \sum_{u \in \mathcal{U}} \pi(u) \Gamma_u^*(\xi) \right\}. \quad (27)$$

The argument proceeds as follows. We formalize time scale separation $\tau_2 \ll \tau_3$ by taking the limit of small

gains $\alpha_r = \varepsilon a_r$, $\varepsilon \rightarrow 0$. Write then $\xi_r^\varepsilon(t) = \xi_r(t/\varepsilon)$. The dynamics (21) entail that

$$\begin{aligned} \xi_r^\varepsilon(t+h) &= \xi_r^\varepsilon(t) + \\ &+ a_r \int_t^{t+h} [U_r'^{-1}(\xi_r^\varepsilon(s)) - x_r(\xi^\varepsilon(s), u(s/\varepsilon))] ds, \end{aligned}$$

where $x(\xi, u)$ is any solution to the conditions (20). Provided the time derivative of ξ_r^ε is uniformly bounded, and if $x_r(\xi, u)$ were continuous in its first argument (a condition that is not met in our present context), we could then invoke the fact that in the interval $[t/\varepsilon, (t+h)/\varepsilon)$, the fraction of time spent in phase u approaches $\pi(u)$. This would then imply that any limiting trajectory of ξ_i^ε must satisfy

$$\begin{aligned} \xi_r(t+h) &= \xi_r(t) + a_r \int_t^{t+h} [U_i'^{-1}(\xi_r(s)) - \\ &- \sum_{u \in \mathcal{U}} \pi(u) x_r(\xi(s), u) + o(h)] ds, \end{aligned}$$

and these are precisely the dynamics we now set out to analyze. General results on the approach we have just sketched are available in [11], under the name of averaging theory.

For future reference, introduce the notation

$$\Gamma^*(\xi) := \sum_{u \in \mathcal{U}} \pi(u) \Gamma_u^*(\xi). \quad (28)$$

The previous lemma allows us to derive the following result.

Theorem 6: Consider the dynamics specified by (21), where $x(t)$ satisfies (27). It then holds that ξ converges to the maximizer of the Lyapunov function

$$L(\xi) := \sum_{r \in \mathcal{F}} \int_0^{\xi_r} U_r'^{-1}(v) dv - \Gamma^*(\xi). \quad (29)$$

Denote the corresponding limiting point by ξ^* , and let $\bar{x}_r := U_r'^{-1}(\xi_r^*)$ denote the corresponding rate. For each $u \in \mathcal{U}$, let $x_r(u)$, $r \in \mathcal{F} \cup \mathcal{M}$ denote a solution to the instantaneous optimization problem (16) with ξ set to ξ^* . Provided these solutions verify the additional consistency relation (25), then the variables $x_r(u)$ solve the maximization problem (24).

Proof: We first identify the convex conjugate function of Γ^* as defined in (28), that we denote by Γ . By Theorem 16.4, p.145 in [16], we find that

$$\Gamma(x) = \inf \left\{ \sum_{u \in \mathcal{U}} G_u(x_u), \sum_{u \in \mathcal{U}} x_u = x \right\},$$

where G_u is the convex conjugate of $\xi \rightarrow \pi(u) \Gamma_u^*(\xi)$. It can easily be verified that this reads $G_u(x) =$

$\pi(u)\Gamma_u(\pi(u)^{-1}x)$. This then yields the following characterization of Γ as

$$\Gamma(x) = \inf \left\{ \sum_{u \in \mathcal{U}} \pi(u)\Gamma_u(x(u)), \sum_{u \in \mathcal{U}} \pi(u)x(u) = x \right\}.$$

By Lemma 1, we then find that the vector \bar{x} maximizes the objective function:

$$\sum_r U_r(\bar{x}_r) - \Gamma(\bar{x}). \quad (30)$$

Moreover, the requirement that the rates $x_r(u)$ each solve the instantaneous optimization problem (16) with $\xi = \xi^*$ implies the following condition:

$$\xi^* \in \partial\Gamma_u(x(u)). \quad (31)$$

On the other hand, let $y_r(u)$ denote rates that solve the optimization problem (24)–(26). It is readily seen, by maximizing over rates $x_r(u)$ with $r \in \mathcal{M}$ first, that the $y_r(u)$, $r \in \mathcal{F}$ are characterized as maximizers of

$$\sum_{r \in \mathcal{F}} U_r(\bar{x}_r) - \sum_{u \in \mathcal{U}} \pi(u)\Gamma_u(x(u)),$$

subject to the relation (25). This admits the equivalent characterization that the average rates \bar{y}_r maximize (30), that the relations (25) holds, and that there exists a vector of multipliers μ_r such that for all $u \in \mathcal{U}$,

$$\mu \in \partial\Gamma_u(y(u)).$$

However, the vectors $x_r(u)$ in the theorem meet these conditions, with the multipliers ξ_r^* . This allows us to conclude. ■

B. Sub-optimal algorithms

The algorithm of the last section is defined for convergence to SYSTEM optimal allocation. By its very nature, the optimal allocation is such that in some phases the rate of a farsighted user drops to zero. Such a radical strategy is not desirable in practice. One issue is that, for networks where charge per unit flow is conveyed through a binary feedback, such as ECN marks or packet losses, users cannot at the same time have zero send rate and still have feedback information about the marginal cost of sending along their path. We propose to circumvent this difficulty in the following way. The idea is for a farsighted user r to use a “local” demand function $g(\cdot/w_r)$ (positive-valued, decreasing), so that when the instantaneous charge per unit flow is p_r , it sends at rate

$$x_r = g(p_r/w_r), \quad (32)$$

where w_r is a parameter adapted by the user at a slower time scale (i.e., τ_3). In the sequel we restrict our attention

to local demand functions derived from α -fair utility functions (15), i.e.

$$g(p) = \left(\frac{1}{p}\right)^{\frac{1}{\alpha}}, \quad \alpha > 0.$$

Clearly, this adds additional constraints to the original optimization problem, and hence we cannot expect to solve it exactly. This has however alleviated the difficulty of having zero send rates, as now users will send at non-zero rates no matter how high the marginal costs p .

Given the constraints (32), which can be achieved for instance by adapting x_r on time scale τ_1 according to

$$\frac{dx_r}{dt} = \kappa_r (w_r x_r^{1-\alpha} - x_r p_r), \quad (33)$$

it remains to specify how w_r should be adapted at time scale τ_3 . We propose two possible schemes for doing that, each making use of the parameter ξ_r , which is as previously adapted according to (21).

The first scheme amounts to taking

$$w_r^{(1)} = U_r'^{-1}(\xi_r) \xi_r \frac{1}{z_r}, \quad (34)$$

where z_r is a variable tracking the average value of $x_r^{1-\alpha}$; for definiteness, we may assume it to be adapted according to

$$\frac{dz_r}{dt} = \alpha'_r (x_r^{1-\alpha} - z_r).$$

The rationale for the choice (34) is as follows. It is such that, given the constraints (32), it will determine the value w_r such that the conservativeness identity (1) is met with equality. We do not try to make a formal claim here, but only give the intuition behind this. Indeed, Equation (21) ensures that limiting points for ξ_r necessarily satisfy

$$\xi_r = U_r'(\bar{x}_r).$$

However, dynamics (33) guarantee that $w_r \overline{x_r^{1-\alpha}} = \bar{x}_r \bar{p}_r$. The desired conservativeness identity follows by plugging expression (34) in there, and using the expression for ξ_r . An approximate version of this scheme, which does not require to maintain the additional variable z_r , is as follows:

$$w_r^{(2)} = [U_r'^{-1}(\xi_r)]^\alpha \xi_r. \quad (35)$$

It may be checked using similar arguments that this eventually enforces the modified conservativeness relation

$$U_r'(\bar{x}_r) \frac{\overline{x_r^{1-\alpha}}}{(\bar{x}_r)^{1-\alpha}} = \bar{p}_r.$$

Remark 1: Both dynamics (34) and (35) correspond to the adaptation of the previous section, as we let $\alpha \rightarrow 0$. If α is positive, but small, we may think of the dynamics of this section as “soft” variants of the dynamics in the earlier section. Also, (34) is expected to be conservative according to (1) with equality, while in view of the previous displayed equation and Jensen’s inequality, (35) is expected to be conservative only when $\alpha \leq 1$. It is noteworthy that for $\alpha = 1$, both adaptations (34) and (35) are the same, and thus the adaptation (35) achieves the conservativeness inequality with equality.

VI. NUMERICAL EXAMPLES

We give a limited set of numerical results for the algorithms introduced in Section V. The results illustrate some properties of equilibria described in Section III. The numerical results are for the model described as follows. We consider a single resource that offers the cost function defined by $C'(x) = 0$, for $0 < x \leq \eta c$, $C'(x) = (\frac{x-\eta c}{c-\eta c})^\gamma$, for $\eta c \leq x < c$, $C'(x) = 1$, $x \geq c$, with the parameters fixed as $\eta = 1/2$, $\gamma = 4$, and the resource capacity $c = 10$. A population of non-persistent myopic users arrive to the resource according to a homogeneous Poisson process with intensity $\lambda = 0.01$ and each arrival is a file transfer with mean size $1/\mu = 20$; this amounts to the load of non-persistent myopic users of $\rho = 1/5$. Two persistent users compete for the resource. All the users have α -fair utility function with the parameters $w_r = 1$ and $\alpha = 2$ (TCP-fairness). Figure 4 shows the competition of two myopic persistent users. Figures 5–7 show the results for one myopic and one farsighted persistent user. The farsighted user in Figure 5 and Figure 6 runs the adaptation of Section V-A, but with different averaging constants. The farsighted user in Figure 7 runs a soft version of the algorithm. All results demonstrate validity of the algorithms; see figure captions for some observations.

VII. CONCLUSION AND FUTURE WORK

In this paper we have introduced a formal framework for specifying network resource allocations under time-diversity. We have introduced the class of so-called farsighted users. We have established a number of characteristic properties of such farsighted users, which distinguish them from their myopic relatives. In particular, such users have a smoothing effect on the network conditions, which we believe is a desirable property. At the same time, such users do obey a conservativeness relation.

We have developed two classes of congestion control algorithms for farsighted users. The first class solves the target SYSTEM optimization problem exactly under a time scale separation assumption. The second class

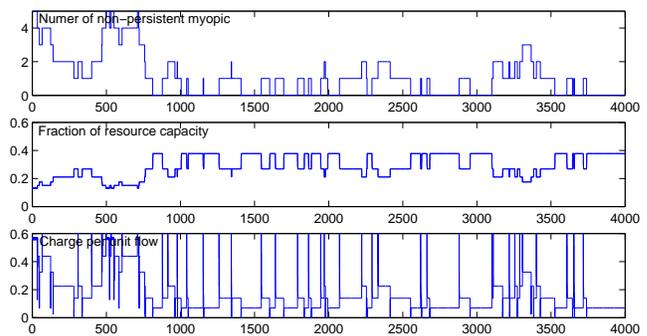


Fig. 4. Two persistent myopic users compete with non-persistent myopic users for a resource. The top graph shows the number of non-persistent myopic users. The middle graph shows allocations for the two persistent users; in this case, the two curves overlap. The bottom graph is charge per unit flow.

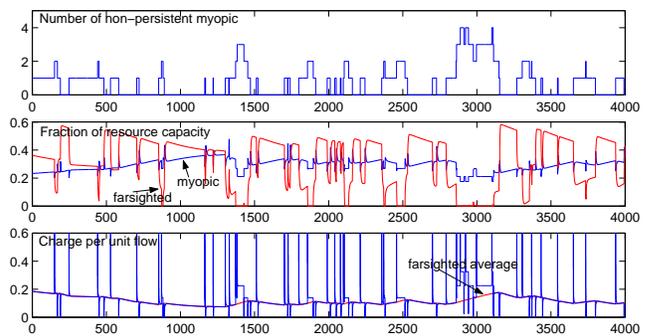


Fig. 5. Same as in Figure 4, but one persistent user is now farsighted. The farsighted user uses the gain parameter $\alpha_r = 0.0001$. The middle graph indicates that the farsighted user tends to use more resources when a few non-persistent users are present in the system, than the competing myopic persistent user. When a larger number of non-persistent users are in the system, the farsighted user “backs-off” by dropping its rate to zero. The farsighted user receives slightly less than 5.5% larger time-average allocation than the persistent myopic user.

solves an approximate version of the original optimization problem, but is more appealing from an implementation point of view.

There are several directions to explore further. One is to explore further properties and establish some found in this paper to hold more generally. Throughout the paper we assumed users are price takers, which makes room to explore farsighted users that anticipate their effect on the prices [15]. Some of the properties may still remain, as with the price taking assumption, after appropriate adjustments; for instance the price equalization property carries over, but it is the first derivative of the charge per unit time with respect to the allocation that is equalized. Another research direction is to investigate the cases when π is not assumed to be fixed in solving the optimization problem. More thorough analysis of the algorithms is needed along with their discrete-event implementations. We anticipate that the

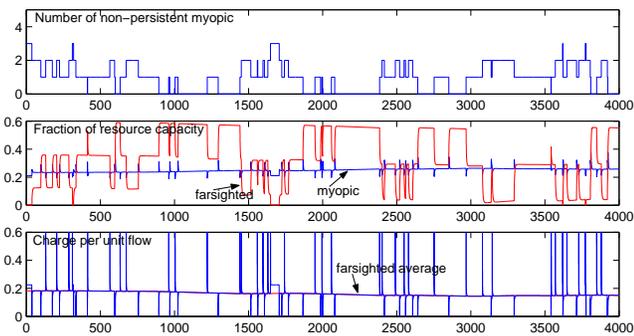


Fig. 6. Same as in Figure 5, but the farsighted user uses the gain parameter $\kappa_r = 0.00001$. Due to larger averaging the allocation (middle graph) is now closer to the equilibrium predicted by the optimization. The farsighted user now receives time-average resource allocation that is slightly more than 30% than that of the competing myopic user. Note that the resource allocation of the persistent myopic user is considerably smoother than in the competition of all myopic users (Figure 4); this we expect by the price equalization.

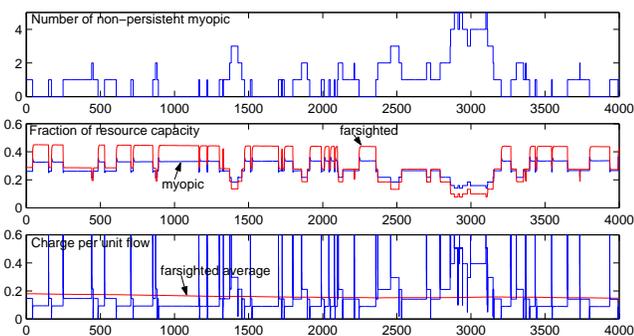


Fig. 7. Same as in Figure 5, but the farsighted user now runs soft version of the adaptation (33), with $\alpha = 1$. As expected, the farsighted user now has no sharp drops to zero; the farsighted user receives slightly more than 16.5% time-average allocation than the competing myopic user.

algorithms proposed in this paper would be amenable to implementation by appropriate modifications of an existing transport protocol such as TCP.

ACKNOWLEDGMENTS

M. V. acknowledges early discussions with Jean-Yves Le Boudec on farsighted species.

REFERENCES

- [1] Dimitri Bertsekas and Robert Gallager. *Data Networks*.
- [2] A. Kuzmanović and E. Knightly. TCP-LP: A Distributed Algorithm for Low Priority Data Transfer. In *Proc. of IEEE INFOCOM 2003*, 2003.
- [3] Sally Floyd and Kevin Fall. Promoting the Use of End-to-End Congestion Control in the Internet. *IEEE/ACM Trans. on Networking*, 7(4):458–472, August 1999.
- [4] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
- [5] R. J. Gibbens and F. P. Kelly. On packet marking at priority queues. *IEEE Trans. on Automatic Control*, 47:1016–1020, 2002.

- [6] Mark Handley, Jitendra Padhye, Sally Floyd, and Jörg Widmer. TCP Friendly Rate Control (TFRC) Protocol Specification, IETF INTERNET-DRAFT, January 2003. RFC 3448, <ftp://ftp.isi.edu/in-notes/rfc3448.txt>.
- [7] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997. a corrected version available from <http://www.statslab.cam.ac.uk/~frank>.
- [8] F. P. Kelly. Mathematical modelling of the Internet. In *Forth International Congress on Industrial and Applied Mathematics*, Edinburgh, Scotland, July 1999. (also available from <http://www.statslab.cam.ac.uk/~frank>).
- [9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49, 1998. (also available from <http://www.statslab.cam.ac.uk/~frank>).
- [10] Peter Key, Laurent Massoulié, and Bing Wang. Emulating Low-Priority Transport at the Application Layer: A Background Transfer Service. In *Proc. of ACM Sigmetrics 2004*, pages 118–129, 2004.
- [11] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2002.
- [12] Srisankar Kunniyur and R. Srikant. End-to-End Congestion Control Schemes: Utility Functions, Random Losses and ECN Marks. In *Proc. of the IEEE INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [13] Anukool Lakhina, Konstantina Papagiannaki, Mark Crovella, Christophe Diot, Eric D. Kolaczyk, and Nina Taft. Structural analysis of network traffic flows. In *Proc. of ACM Sigmetrics'04*, June 2004.
- [14] Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *submitted to IEEE/ACM Trans. on Networking*, 1999.
- [15] Johari R. and Tsitsiklis J. N. Efficiency Loss in Network Resource Allocation Game. *To Appear in Mathematics of Operations Research*, 2004. <http://web.mit.edu/rjohari/www/>.
- [16] Rockafellar R. T. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [17] A. Venkataramani, R. Kokku, and M. Dahlin. TCP Nice: A Mechanism for Background Transfers. In *Proc. of Operating Systems Design and Implementation*, December 2002.
- [18] Milan Vojnović and Jean-Yves Le Boudec. On the Long-Run Behavior of Equation-Based Rate Control. *To Appear in IEEE/ACM Trans. on Networking*, April 2005.
- [19] Milan Vojnović, Jean-Yves Le Boudec, and Catherine Boutremans. Global Fairness of Additive-Increase and Multiplicative-Decrease with Heterogeneous Round-Trip Times. In *Proc. of IEEE Infocom 2000*, Tel-Aviv, Israel, March 2000.

APPENDIX

A. Proof of Proposition 1

Consider average throughput of a myopic user:

$$\bar{x}_m(k) = \sum_{u \in \mathcal{U}: x_f(u) > 0} \pi(u) \bar{x}_f(k) + \sum_{u \in \mathcal{U}: x_f(u) = 0} \pi(u) x_m(u).$$

Differentiating with respect to k , we obtain

$$\frac{\partial \bar{x}_m(k)}{\partial k} = \left(\sum_{u \in \mathcal{U}: x_f(u) > 0} \pi(u) \right) \frac{\partial \bar{x}_f(k)}{\partial k} + \sum_{u \in \mathcal{U}: x_f(u) = 0} \pi(u) \frac{\partial x_m(u)}{\partial k}. \quad (36)$$

The term $\partial \bar{x}_f(k) / \partial k$ is already derived in (6). It remains to determine $\partial x_m(u) / \partial k$, for u such that $x_f(u) = 0$. Note that for each such u , $x_m(u)$ is the unique solution of $(n -$

$k)x_m(u) = g_u(x_m(u))$. Differentiating the last expression with respect to k , we obtain

$$(n-k)\frac{\partial x_m(u)}{\partial k} - x_m(u) = g'_u(x_m(u))\frac{\partial x_m(u)}{\partial k}.$$

Denote the function $h_u(x) = g_u(x) - g'_u(x)x$, $x \geq 0$.² Substituting $\partial \bar{x}_f(k)/\partial k$ in (6) and $\partial x_m(u)/\partial k$ in the last display into (36), we obtain

$$\frac{\partial \bar{x}_m(k)}{\partial k} = \left(\sum_{u \in \mathcal{U}: x_f(u)=0} \pi(u) \right) \left(-\frac{\bar{x}_f(k)^2}{\sum_{u \in \mathcal{U}: x_f(u)>0} \tilde{\pi}(u)h_u(\bar{x}_f(k))} + \sum_{u \in \mathcal{U}: x_f(u)=0} \tilde{\pi}^0(u) \frac{x_m(u)^2}{h_u(x_m(u))} \right)$$

where $\tilde{\pi}(u) = \pi(u)/\sum_{v \in \mathcal{U}: x_f(v)>0} \pi(v)$ and $\tilde{\pi}^0(u) = \pi(u)/\sum_{v \in \mathcal{U}: x_f(v)=0} \pi(v)$. In view of the last above display, $\partial \bar{x}_m(u)/\partial k < 0$ is equivalent to saying

$$\frac{\bar{x}_f(k)^2}{\sum_{u \in \mathcal{U}: x_f(u)>0} \tilde{\pi}(u)h_u(\bar{x}_f(k))} > \sum_{u \in \mathcal{U}: x_f(u)=0} \tilde{\pi}^0(u) \frac{x_m(u)^2}{h_u(x_m(u))}. \quad (37)$$

In the sequel, we show that there exist cases when the last inequality holds. First note that, from (5), the average rate $\bar{x}_f(k)$ of a farsighted user is given by

$$k\bar{x}_f(k) = \sum_{u \in \mathcal{U}} \pi(u) [g_u(\bar{x}_f(k)) - (n-k)\bar{x}_f(k)]^+.$$

A phase u is bad, i.e. $x_f(u) = 0$, if $g_u(\bar{x}_f(k)) \leq (n-k)\bar{x}_f(k)$.

Let there be only two phases, $\mathcal{U} = \{0, 1\}$. Fix $\pi(0) \in (0, 1)$ and a decreasing positive-valued function $g_1(\cdot)$. Let $\bar{x}_f(k)$ be given by $k\bar{x}_f(k) = (1 - \pi(0))(g_1(\bar{x}_f(k)) - (n-k)\bar{x}_f(k))$. Let $g_0(\cdot)$ be any decreasing function such that $g_0(\bar{x}_f(k)) \leq (n-k)\bar{x}_f(k)$. Then, $u = 0$ is bad phase and 1 is good phase, and $\bar{x}_f(k)$ is indeed average rate of a farsighted user. In view of (37), $\partial \bar{x}_m(k)/\partial k < 0$ is equivalent to

$$\frac{\bar{x}_f(k)^2}{h_1(\bar{x}_f(k))} > \frac{x_m(0)^2}{h_0(x_m(0))}.$$

The left side of the inequality is fixed. The right side of the inequality can be made arbitrarily small by choosing appropriate function $g_0(\cdot)$. Hence, it can be $\partial \bar{x}_m(k)/\partial k < 0$. The proposition is proved.

B. Proof of Theorem 4

The proof is to solve an instance of SYSTEM(U, C), which in the prevailing setting can be rewritten as

$$\begin{aligned} & \text{maximize} && U_{\mathcal{F}} \left(\sum_{u \in \mathcal{U}} \pi(u)(c - ux_{\mathcal{M}}(u)) \right) + \\ & && + \sum_{u \in \mathcal{U}} \pi(u)uU_{\mathcal{M}}(x_{\mathcal{M}}(u)) \\ & \text{over} && 0 \leq ux_{\mathcal{M}}(u) \leq c, \quad u \in \mathcal{U}. \end{aligned}$$

²For a given $x > 0$, $h_u(x)$ is the value at which the tangent on the function $g_u(\cdot)$ at x , intersects the ordinate.

The Lagrangian is, for $\mu_u \geq 0$,

$$\begin{aligned} L(x_{\mathcal{M}}, \mu) = & U_{\mathcal{F}} \left(\sum_{u \in \mathcal{U}} \pi(u)(c - ux_{\mathcal{M}}(u)) \right) + \\ & + \sum_{u \in \mathcal{U}} \pi(u)uU_{\mathcal{M}}(x_{\mathcal{M}}(u)) + \\ & + \sum_{u \in \mathcal{U}} \mu_u(c - ux_{\mathcal{M}}(u)). \end{aligned}$$

An optimum satisfies the first order optimality conditions

$$U'_{\mathcal{M}}(x_{\mathcal{M}}(u)) = U'_{\mathcal{F}} \left(\sum_{n \in \mathcal{U}} \pi(n)(c - nx_{\mathcal{M}}(n)) \right) + \mu_u, \quad (38)$$

for all $u \in \mathcal{U} - \{0\}$, along with the complementarity relations

$$\mu_u(c - ux_{\mathcal{M}}(u)) = 0, \quad \mu_u \geq 0, \quad u \in \mathcal{U}.$$

Now, for $u \in \mathcal{U}$ such that $ux_{\mathcal{M}}(u) < c$, we have $\mu_u = 0$; for such an u the right-hand side in (38) does not depend on u , and thus $x_{\mathcal{M}}(u) = x^*$, where x^* is the solution of (7). Else, for $u \in \mathcal{U} - \{0\}$ such that $ux_{\mathcal{M}}(u) = c$, indeed, $x_{\mathcal{M}}(u) = c/u$. The assertion is proved.

C. Proof of Theorem 5

We first show that condition (LP) is necessary for farsighted users to be low-priority users, for arbitrary convex non-decreasing cost functions C . Indeed, given some fixed $x \geq 0$, select such a cost function C satisfying $U'_{\mathcal{M}}(x) = C'(x)$. Recall that u^* is the threshold on number of myopic users defining good phases for farsighted users, so that they are low-priority if and only if $u^* = 0$. Note that $u^* = 0$ is equivalent to the following set of conditions:

$$\begin{aligned} U'_{\mathcal{F}}(\pi(0)x_{\mathcal{F}}(0)) &= C'(fx_{\mathcal{F}}(0)), \\ U'_{\mathcal{F}}(\pi(0)x_{\mathcal{F}}(0)) &\leq C'(x_{\mathcal{M}}(1)), \\ U'_{\mathcal{M}}(x_{\mathcal{M}}(1)) &= C'(x_{\mathcal{M}}(1)). \end{aligned} \quad (39)$$

The third relation in (39) is met iff $x_{\mathcal{M}}(1) = x$, by strict concavity of the utility function $U_{\mathcal{M}}$. The first two relations imply that $x = x_{\mathcal{M}}(1) \geq fx_{\mathcal{F}}(0)$. By concavity of $U_{\mathcal{F}}$, we then have

$$\begin{aligned} U'_{\mathcal{F}} \left(\frac{\pi(0)x}{f} \right) &\leq U'_{\mathcal{F}}(\pi(0)x_{\mathcal{F}}(0)) \\ &\leq U'_{\mathcal{M}}(x), \end{aligned}$$

where the last inequality follows from (39). Since x was arbitrary, this establishes the claimed necessary condition.

We now prove the converse implication. Suppose (LP) holds. Let C be any convex cost function. We now establish that the corresponding relations (39) hold. Let $x_{\mathcal{M}}(1)$ denote the solution to the third relation of (39), i.e. $U'_{\mathcal{M}}(x_{\mathcal{M}}(1)) = C'(x_{\mathcal{M}}(1))$.

By assumption (LP), we deduce that

$$U'_{\mathcal{F}} \left(\frac{\pi(0)x_{\mathcal{M}}(1)}{f} \right) \leq C'(x_{\mathcal{M}}(1)).$$

By convexity of C and concavity of $U_{\mathcal{M}}$, this implies that there exists a solution $x_{\mathcal{F}}(0)$ to the first two relations in (39).