

# Achieving network stability and user fairness through admission control of TCP connections

Andrés Ferragut and Fernando Paganini  
Telecommunications Engineering, Universidad ORT,  
Montevideo, Uruguay

**Abstract**—This paper studies a network under TCP congestion control, in which the number of flows per user is explicitly taken into account. We present a control law for this variable that, in combination with congestion control, induces as equilibrium the maximization of social welfare measured at the level of users, rather than the level of TCP connections. We use fluid flow models to prove stability theorems on the dynamics of the overall system, combining the dynamics of flows with the dynamic rates and prices of congestion control. We then develop an admission control policy for discrete TCP flows, that emulates the continuous behavior, and is modeled as a Markov chain. We present simulation studies of the overall system, which exhibit its stability and the desired user-level fairness behavior.

## I. INTRODUCTION

A fundamental problem in networking concerns the distribution of network resources efficiently and fairly. Since the seminal work of Kelly et al. [6], it has become customary to frame such questions in the language of welfare economics, postulating that resources should be allocated so as to solve a *network utility maximization* problem. Recent surveys to the now abundant literature on this topic are [3], [13].

In particular, these models apply to Internet congestion control [9], [13], by associating to each TCP flow a utility function that determines their response in rate to congestion signals or “prices”. The resulting equilibrium, if reached, achieves a notion of flow level fairness. In the language of [6], the above is the “NETWORK problem”, on top of which it was proposed to add a “USER problem” through which users could express their preferences through a choice of weight in the TCP utility function, making the overall equilibrium optimize overall welfare, the “SYSTEM problem”. What has been lacking in both literature and practice is an implementation of this outer loop by the users. Without it, one is left with a network that strives to impose fairness between individual TCP flows, not user level fairness. In particular, as recently argued eloquently by Briscoe [2], it is useless to make flow rates fair if users are allowed to open an unspecified amount of such flows. Indeed, p2p applications in the Internet are known to use precisely this mechanism to push for a larger share of the bandwidth “pie”.

Failure to control the number of TCP flows not only compromises fairness, but it can also jeopardize network stability. If many users resorted to such greedy tactics, the end result could well be a “tragedy of the commons” scenario with large

numbers of connections, each carrying minuscule amounts of traffic. This issue relates to results [1], [4] on stochastic stability of the number of flows in a network, when these arrive randomly and are served according to some flow level fairness (max-min fairness, or  $\alpha$ -fairness, [12]). The stability condition is that the overall exogenous user rate demand must fit within network capacity, a condition not affected by the rate control of TCP: a user with high demand who is getting a low rate from TCP, will open as many TCP connections as needed to obtain its desired rate. Furthermore, if users happen to desire an overall demand that cannot be met, stability fails and the number of TCP connections will grow without bounds. The end result is a loss of performance not just for the greedy users, but for the entire network, as also discussed in [5].

Motivated by these considerations, in this paper we propose to recover the user-level efficiency and fairness model of the SYSTEM problem, by controlling the number of TCP connections admitted in the network. In other words, let the user weight parameter be the number of TCP connections it opens; we propose to control this parameter in feedback so as to achieve an equilibrium that maximizes user welfare. In Section II we study this type of mechanism at the level of fluid models, as if the number of flows were a continuous variable. We propose a control law that adapts this variable based on current congestion prices, and the per-flow rate obtained by TCP; the equilibrium of the overall system is designed to solve the SYSTEM problem, and we obtain analytical results on dynamic stability of this equilibrium.

In Section III we move closer to implementation by considering the discrete nature of flows. We emulate the fluid flow law with an admission control policy, defined in terms of a user demand curve that reflects a service level agreement between the user and the network. This admission control stabilizes the Markov chain that results from stochastic flow arrivals, for an arbitrary exogenous load. In Section IV we develop a simulator for this system with two timescales: stochastic arrivals and workloads, and congestion control via differential equations. We present simulations that exhibit the features of the control algorithm. Conclusions are given in Section V.

## II. USER-LEVEL FAIRNESS CONTROL FOR MULTIPLE TCP FLOWS

We consider a network composed of links, indexed by  $l$ , with capacity  $c_l$ , and a set of end-to-end users, indexed by  $i$ , which send their flow through a single path, characterized by

{ferragut,paganini}@ort.edu.uy. Research supported in part by PDT-Uruguay, and by AFOSR-US, grant FA9550-06-1-0511.

the routing matrix  $R$ , with entries

$$R_{li} = \begin{cases} 1 & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}.$$

Following [6], assign each user a utility  $U_i(\rho_i)$ , increasing and concave function of the *aggregate rate*  $\rho_i$  it obtains from the network, potentially through multiple TCP flows. Specifically  $\rho_i = n_i x_i$ , where  $x_i$  is the rate assigned by TCP to a single flow, and  $n_i$  is the number of flows. The aggregate rate in link  $l$  from all sources is

$$y_l = \sum_i R_{li} \rho_i. \quad (1)$$

In matrix form, we write  $y = R\rho$ . We are now in a position to state the desired resource allocation:

**Problem 1 (SYSTEM):** Maximize  $\sum_i U_i(\rho_i)$ , subject to link capacity constraints  $R\rho \leq c$ .

This is the same problem considered by [6]; its rationality as a model for fairness is that users are considered as an entity, in terms of the overall bandwidth resources they consume. The utility functions in Problem 1 are a degree of freedom that may reflect different models of fairness or service differentiation. A useful family in this regard is the well-known  $\alpha$ -fairness model from [12],

$$U_i(\rho) = k_i \frac{\rho^{1-\alpha}}{(1-\alpha)}, \text{ with } \alpha > 0. \quad (2)$$

Each TCP flow of source  $i$  obtains a rate  $x_i$  that is determined in feedback, based on congestion price  $q_i$  received from the network, following a control law determined by the underlying TCP protocol. A simple fluid model of this relationship is the demand curve

$$x_i = f_{TCPi}(q_i); \quad (3)$$

equivalently, TCP is maximizing  $U_{TCPi}(x_i) - q_i x_i$  for an increasing, concave utility function  $U_{TCPi}$  satisfying  $U'_{TCPi} = f_{TCPi}^{-1}$ . Another model of rate control, with the same equilibrium behavior, is the ‘‘primal’’ law

$$\dot{x}_i = \kappa(x_i)[U'_{TCPi}(x_i) - q_i]. \quad (4)$$

For more details see e.g. [13]. We emphasize the distinction between the utility function  $U_{TCPi}$  that models the *protocol behavior* at a fast time scale, and the utility  $U_i$  that models the *user demand*, at a slower time-scale. For instance, if the network runs TCP-Reno, its demand curve can be approximated by the square-root formula of [11], corresponding to  $U_{TCP}$  of the  $\alpha$ -fairness kind with  $\alpha = 2$ . This utility need not reflect the demand of the user behind this flow.

Concentrating on the fast time-scale of TCP congestion control where the number of flows  $n_i$  remains constant, the resource allocation can be characterized by another network utility maximization problem:

**Problem 2 (TCP Congestion control):**

Maximize  $\sum_i n_i U_{TCPi}(x_i)$ , subject to link capacity constraints  $\sum_i R_{li} n_i x_i \leq c_l$  for each  $l$ .

Here, as in the standard formulation (see [13], Chapter 9), each TCP flow is assigned a utility  $U_{TCPi}$ , and the congestion control portion maximizes its sum; the objective above groups

terms for the  $n_i$  flows of each user. We can also call the above a ‘‘NETWORK’’ problem as the one considered in [6] with weighted log utilities; here, the weights are given by the numbers of flows  $n_i$ .

TCP congestion control can be related with Problem 2 if one assumes link congestion prices  $p_l$  generated by

$$\dot{p}_l = \gamma_l (y_l - c_l)_{p_l}^+ := \begin{cases} y_l - c_l, & \text{if } y_l > c_l \text{ or } p_l > 0; \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

and source prices

$$q_i = \sum_l R_{li} p_l, \quad (6)$$

in vector form  $q = R^T p$ . Then the equilibrium rates solve Problem 2, and are globally attractive [13].

This brings us to our main problem. For each fixed set of  $n_i$ 's, the network returns a certain set of TCP rates  $x_i$  per flow, hence an overall rate  $\rho_i = n_i x_i$  per user. It is not difficult to see that the mapping  $n_i \mapsto \rho_i$  for one user in equilibrium is nondecreasing, for all other  $n_j$ 's fixed. Therefore each individual user, acting selfishly, has an incentive to increase its number of connections; if one user (or a subset) behaves in this manner while the rest keep  $n_i$  unchanged, the latter group would be pushed out of the network as their  $x_i$ 's diminish. The scenario where everybody greedily increases  $n_i$  without bound could hypothetically give fair  $\rho_i$ 's, but at a complexity cost which is clearly also undesirable.

In conclusion, if one seeks to control the resource allocation at the user level, the numbers of flows  $n_i$  should be controlled. The purpose of this section is to find a control law for  $n_i$  at a slow time-scale, that in conjunction with TCP at a faster time-scale, reaches user rates  $\rho_i$  compatible with Problem 1. We now propose such an algorithm at the level of fluid flow models, and study its analytical properties.

### Control of the number of flows:

$$\dot{n}_i = \beta (U_i'^{-1}(q_i) - \rho_i) \quad (7)$$

where  $\beta$  is a positive constant. The above equation treats  $n_i$  as a continuous variable, later on we will consider more practical versions with discrete numbers of flows.

The intuition behind this control law is straightforward: the right hand-side compares the user's demand at the current congestion price, with the rate obtained currently from the network. This difference dictates whether the number of connections should be increased, or decreased. The various layers of the user are collectively described by

$$\dot{n}_i = \beta (U_i'^{-1}(q_i) - \rho_i); \quad (8a)$$

$$x_i = f_{TCP}(q_i); \quad (8b)$$

$$\rho_i = n_i x_i. \quad (8c)$$

**Remark:** Alternatively, (8b) could be replaced by the primal law (4).

Combining this user law with equations (1) and (6) that characterize the network, and (5) that characterizes prices of network links, we have a complete description of the control loop. This overall model is analogous to the congestion control picture of e.g. [9], but with the overall rates of users as variables, instead of the rates of individual flows. Since this

model is directed to the slower time scale of the birth and death of flows, we will not consider network delays in the dynamics.

An equilibrium point  $(n^*, x^*, \rho^*, y^*, p^*, q^*)$  of the overall system will satisfy  $y^* = R\rho^*$ ,  $q^* = R^T p^*$ ,  $\rho_i^* = n_i^* x_i^*$  and

$$U'_i(\rho_i^*) = q_i^*, \quad (9)$$

$$p_l^*(c_l - y_l^*) = 0, \quad (10)$$

$$U'_{TCPi}(x_i^*) = q_i^*. \quad (11)$$

(9-10) imply the Karush-Kuhn-Tucker conditions for optimality in Problem 1 (see [13], replacing  $x$  with  $\rho$ ). Also, (10-11) imply the KKT conditions for Problem 2. This follows from considering the corresponding Lagrangian. Therefore, the equilibrium prices serve simultaneously as Lagrange multipliers for both optimization problems. This justifies the use of the same price variable for TCP congestion control and admission control.

#### A. Local stability for the network case

We wish to establish the stability of the equilibrium point under the proposed dynamics. In this section, we provide a local result following a *passivity approach* (cf. [7]), that was first introduced in the congestion control context by Wen and Arcak [14]. In that paper, the authors show that the system  $(\rho - \rho^*) \mapsto (q - q^*)$  with links applying the dual law (5) is passive with storage function  $V_{net}(p) = \frac{1}{2\gamma} \sum_{i=1}^L (p_i - p_i^*)^2$ . In other words, they establish the dissipation inequality

$$\dot{V}_{net} \leq \sum_i (q_i - q_i^*)(\rho_i - \rho_i^*) \quad (12)$$

along trajectories of the system. To establish stability of the entire closed loop, it suffices to show that the user system  $-(q_i - q_i^*) \mapsto (\rho_i - \rho_i^*)$  is itself passive for each  $i$ . Indeed, this would mean that there is a storage function  $V_i(n)$  for the subsystem of user  $i$ , satisfying

$$\dot{V}_i \leq -(q_i - q_i^*)(\rho_i - \rho_i^*) \quad (13)$$

Adding (12) and (13) for each  $i$  implies that  $V := V_{net} + \sum_i V_i$  is a Lyapunov function for the closed loop system, from which stability follows. Also, a slight refinement with strict inequalities in (13) (strict passivity) implies asymptotic stability. At the time of writing we are only able to establish this passivity *locally* around equilibrium, through the corresponding linearization. For this we resort to a characterization of linear passivity in terms of the system transfer function  $H(s)$  (see [7]): a system is passive if it is stable and  $\text{Re}(H(j\omega)) \geq 0$  for all  $\omega \in \mathbb{R}$ , and strictly passive if the above inequality is strict.

**Lemma 1:** The linearization of the system  $-(q_i - q_i^*) \mapsto (\rho_i - \rho_i^*)$  defined by equations (8) has a strictly passive transfer function  $H_i(s)$ .

*Proof:* We omit for brevity the derivation of the linearization and its corresponding Laplace transform. The end result is (dropping the subindex  $i$ ):

$$H(s) = \frac{n^* b s + \beta x^* a}{s + \beta x^*}, \quad (14)$$

$$\text{where } a = -\frac{\partial U'^{-1}(q)}{\partial q}(q^*), \quad b = -\frac{\partial f_{TCP}(q)}{\partial q}(q^*),$$

positive due to the concavity of utility. This is a system of the “lead-lag” type. It is straightforward to see the the curve  $H(j\omega)$  (the Nyquist plot) is a circle on the right half-plane, through the points  $H(0) = a$  and  $H(\infty) = n^* b$ . Therefore  $\text{Re}(H(j\omega)) > 0$  as was to be proved. ■

**REMARK:** We could also replace (8b) with (4) as remarked before. In that case, the transfer function is

$$H(s) = \frac{(kn^* + \beta x^* a)s + (k\beta x^* ab)}{(s + kb)(s + \beta x^*)}$$

which can be shown to be passive as well.

We can now state the main result of this section:

**Theorem 2:** The equilibrium of the system given by (8), (1), (6) and (5) is asymptotically stable.

*Proof sketch:* Given the strict passivity of the linearized system  $-\delta q_i \rightarrow \delta \rho_i$ , there exists (see [7]) a quadratic storage function satisfying a local dissipation inequality of the form (13). The construction of the global Lyapunov function follows as described above. ■

#### B. Global stability for the single link case

Although the above result is only local, simulations of the differential equations indicate that stability holds globally as well. In this section, we give a global stability proof under more restrictive conditions:

- We consider a network with a single bottleneck of capacity  $C$  which is shared by  $N$  users, each with utility function  $U_i$ ,  $i = 1, \dots, N$
- We assume the congestion control algorithm of all users follows the same demand curve  $f_{TCP}$ . For instance, this could correspond to everyone running TCP-Reno with the same round-trip time.
- We adopt a separation of time-scales argument. For fixed numbers of flows  $n = (n_1, \dots, n_N)$ , we assume the rates and prices of the congestion control loop converge quickly to the solution of Problem 2. Denote by  $\hat{x}_i(n)$ , the corresponding TCP rate, by  $\hat{q}(n) = \hat{p}(n)$  the congestion price of the link.  $\hat{\rho}_i(n) = n_i \hat{x}(n)$  is the  $i$ -th user's aggregate rate. The hat indicates the equilibrium values for the congestion control at the fast timescale. We then model the slow dynamics of the  $n_i$  variables by

$$\dot{n}_i = \beta [U_i'^{-1}(\hat{q}(n)) - \hat{\rho}_i(n)]. \quad (15)$$

Before stating the stability theorem, note that since there is a single price and a common TCP demand curve, the individual flow rates are common to all users:

$$\hat{x}_i(n) = \frac{C}{\sum_{i=1}^N n_i} =: \hat{x}(n) \quad \forall i. \quad (16)$$

From here we have  $\hat{q}(n) = U'_{TCP}(\hat{x}(n))$  and  $\hat{\rho}_i(n) = n_i \hat{x}(n)$  which substituted in (15) give an autonomous system of differential equations in the state  $n$ .

**Theorem 3:** The solution of the differential equations in (15) converges globally to an equilibrium point  $n^*$  such that  $\rho^* = \hat{\rho}(n^*)$  is the solution of Problem 1.

*Proof:* First note that an equilibrium of (15), together with (16), automatically ensures the KKT conditions (9-10) which characterize the optimum of Problem 1. We thus focus on the global stability proof.

Denote by  $|n| = \sum_{i=1}^N n_i$  the total number of connections. This way  $\hat{x}(n) = C/|n|$ . Adding equations (15) in  $i$  we obtain:

$$\dot{|n|} = \beta \left[ \sum_{i=1}^N U_i'^{-1}(\hat{q}(n)) - \sum_{i=1}^N \hat{\rho}_i(n) \right]$$

Observing that  $\sum_{i=1}^N \hat{\rho}_i(n) = C$  and that

$$\hat{q}(n) = U_{TCP}'(\hat{x}(n)) = U_{TCP}'(C/|n|)$$

we can rewrite the dynamics of  $|n|$  as:

$$\dot{|n|} = \beta \left[ \sum_{i=1}^N f_i \left( \frac{C}{|n|} \right) - C \right], \quad (17)$$

where we introduce the notation  $f_i(x) = U_i'^{-1}(U_{TCP}'(x))$ . The function  $f_i$  is increasing since it is the composition of two decreasing functions. The preceding equation is an autonomous differential equation in the scalar variable  $|n|$ .

Denoting by  $|n^*|$  the equilibrium point of (17), define the following candidate Lyapunov function:

$$V_1(|n|) = \frac{1}{2\beta} (|n| - |n^*|)^2 \quad (18)$$

By differentiation we obtain:

$$\begin{aligned} \dot{V}_1 &= (|n| - |n^*|) \left[ \sum_{i=1}^N f_i \left( \frac{C}{|n|} \right) - C \right] \\ &= (|n| - |n^*|) \left[ \sum_{i=1}^N f_i \left( \frac{C}{|n|} \right) - f_i \left( \frac{C}{|n^*|} \right) \right], \end{aligned}$$

where the last step uses the equilibrium condition. Noting that each  $f_i \left( \frac{C}{|n|} \right)$  is a decreasing function of  $|n|$ , we have that  $\dot{V}_1 \leq 0$ , the inequality being strict away from equilibrium. This shows that  $|n|$  converges to the equilibrium  $|n^*|$  and consequently, the rate of each connection  $\hat{x}(n) \rightarrow x^* = C/|n^*|$ .

We return now to equation (15). By defining  $\delta n_i = n_i - n_i^*$  we can rewrite it as:

$$\begin{aligned} \dot{\delta n}_i &= \beta [f_i(\hat{x}(n)) - n_i \hat{x}(n)] \\ &= -\beta \hat{x}(n) \delta n_i + \beta [f_i(\hat{x}(n)) - n_i^* \hat{x}(n)] \end{aligned}$$

The term  $f_i(\hat{x}(n)) - n_i^* \hat{x}(n) = b(t)$  vanishes as  $t \rightarrow \infty$  since  $\hat{x}(n) \rightarrow x^*$  and  $f_i(x^*) = \rho_i^* = n_i^* x^*$ .

Now take the Lyapunov function  $V_2(\delta n_i) = \frac{1}{2\beta} (\delta n_i)^2$ , its derivative is:

$$\dot{V}_2 = -\hat{x}(n) (\delta n_i)^2 + b(t) \delta n_i \quad (19)$$

Let  $\varepsilon > 0$  be arbitrary, and choose  $t_0$  such that  $\hat{x}(t) > \frac{x^*}{2}$  and  $|b(t)| < \frac{x^*}{2} \varepsilon$  whenever  $t > t_0$ . Then, in the region where  $|\delta n_i| \geq \varepsilon$  and for  $t > t_0$  we have that:

$$|b(t) \delta n_i| < \frac{x^*}{2} \varepsilon |\delta n_i| < \hat{x}(n) \varepsilon |\delta n_i| < \hat{x}(n) (\delta n_i)^2$$

and this inequality proves that in equation (19), we have  $\dot{V}_2 < 0$  for  $|\delta n_i| \geq \varepsilon$  and  $t > t_0$ . This shows that eventually,  $\delta n_i$  will reach the set  $|\delta n_i| < \varepsilon$  and since this is true for arbitrary  $\varepsilon$ , we have that  $\delta n_i \rightarrow 0$ . ■

### III. FAIR ADMISSION CONTROL:

#### STOCHASTIC MODEL AND IMPLEMENTATION.

In this section we expand our control models to take into account the discrete nature of connections. A direct discretization of the fluid law (7) would involve both starting and terminating flows, but the latter operation is disruptive in practice. For this reason, in this section we consider a control policy that restricts the *admission* of new connections only, and relies on their natural termination, assuming each carries a finite workload.

We develop a stochastic model for this process, of a similar nature to the one in [1], [4], which we begin by reviewing. In their model, connections of user  $i$  arrive as a Poisson process of intensity  $\lambda_i$  and each connection carries an exponentially distributed workload of mean  $1/\mu_i$ . The rate at which this workload is served depends on congestion control, resulting in a continuous time Markov chain model for  $n(t) = (n_1(t), \dots, n_N(t))$ , the number of user connections at any given time. The  $Q$ -matrix transition rates are given by:

$$q_{n, n+e_i} = \lambda_i; \quad q_{n, n-e_i} = \mu_i n_i \hat{x}_i(n),$$

where  $e_i$  is  $i$ -th coordinate vector, and  $\hat{x}_i(n)$  is the rate allocated to each connection by TCP, as a function of the number of flows present. [4] considered max-min fairness, and [1] generalized this to Problem 2 with  $\alpha$ -fair utilities as in (2). Other recent work in this direction is [10], and [8] has extensions to wireless scheduling. The main result of [1], [4] is that the necessary and sufficient condition for network stability is:

$$\sum_{i=1}^N R_{li} \frac{\lambda_i}{\mu_i} < c_l \quad \forall l. \quad (20)$$

This says that stability is determined by the exogenous demands  $\frac{\lambda_i}{\mu_i}$ , not the lower level resource allocation of TCP. Users with high demand, if served by a small TCP rate, will accumulate unfinished work and hence increase the number of active connections, as new arrivals are generated. If the overall demand satisfies (20), then an equilibrium in  $n$  is reached, but this could be arbitrarily biased towards these greedy users. And if the condition fails, we have an unstable Markov chain, hence an increasing number of unfinished connections.

Analyzing this situation, it appears that the two main tasks of congestion control (stability and fairness) are not being addressed. Again, this is our motivation for controlling the number of flows; we now propose an admission control mechanism inspired on the results of Section II, to guarantee stability and user fairness.

As in Section II, associate the utility function  $U_i(\rho_i)$  with user  $i$ . The algorithm proceeds as follows: at the ingress node we measure the aggregate load  $\rho_i$  the user obtains from the network, and we compare it with the demand curve  $U_i'^{-1}(q_i)$  for the current network price  $q_i$ . The admission decision for a new connection is:

$$\begin{aligned} \text{If } U_i'^{-1}(q_i) > \rho_i &\rightarrow \text{admit connection.} \\ \text{If } U_i'^{-1}(q_i) \leq \rho_i &\rightarrow \text{discard connection.} \end{aligned} \quad (21)$$

A queueing model similar to the one in [1], [4] can be developed for this system. If  $n(t) = (n_1(t), \dots, n_N(t))$  are the

number of connections of each user, then  $n(t)$  is a continuous time Markov chain with transition rates:

$$q_{n,n+e_i} = \lambda_i \mathbf{1}_{\{U_i'^{-1}(\hat{q}_i(n)) > \hat{p}_i(n)\}} \quad (22a)$$

$$q_{n,n-e_i} = \mu_i n_i \hat{x}_i(n) \quad (22b)$$

where as in Section II-B,  $\hat{x}_i(n)$  and  $\hat{q}_i(n)$  are the connection rates and route prices in equilibrium of the congestion control Problem 2, and  $\hat{p}_i(n) = n_i \hat{x}_i(n)$  is the aggregate rate of user  $i$ .

With this policy, stability is not an issue anymore: under mild restrictions on utilities, the admission control prevents the process  $n(t)$  from escaping towards infinity. We state the following result under the assumptions of Theorem 3. The proof is omitted for brevity.

**Proposition 4:** Consider a single bottleneck of capacity  $C$ , shared by  $N$  users with identical TCP demand curves satisfying  $U_{TCP}'(x) \rightarrow \infty$  when  $x \rightarrow 0$ . Then the trajectories of the Markov chain described by (22) converge to a bounded set  $\mathcal{N}$ .

The remaining issue is establishing the fairness of the proposed admission control. The results of section II suggest that our proposal should enable a fair allocation of resources according to Problem 1. In the next section we show simulations to support this conclusion.

### Remarks on implementation

A denial of admission for a new TCP flow can be performed by the ingress router, by simply discarding TCP packets with the *SYN* flag activated that come from a certain user. The decision would be made following a utility function, which we can regard as a service-level agreement (SLA) between a customer and the network. To enforce this policy, the edge router must know the user aggregate rate and congestion price. Both of these impose some inevitable overhead on the router, at a per-user level.

We can also think of the ingress router of an operator where the “user” is an aggregate of users, served by a common SLA, the proposal can be more attractive, and equilibria with large numbers of TCP flows become more natural.

## IV. SIMULATIONS: STOCHASTIC LOAD WITH ADMISSION CONTROL AND CONGESTION CONTROL

In order to test the admission control, we developed a simulation tool based on JAVA that generates connections for each user following a Poisson process of intensity  $\lambda_i$ . Each arriving connection has an exponentially distributed workload with mean  $1/\mu_i$ . The connections present in the system at any time perform a dual congestion control algorithm, simulated through an Euler discretization of the differential equations. This determines the rate allocated to each connection, and the workload is consumed according to this time-varying rate.

We selected two simulation scenarios based on the linear network shown in Figure 1.




Fig. 1. Linear network used in simulations.

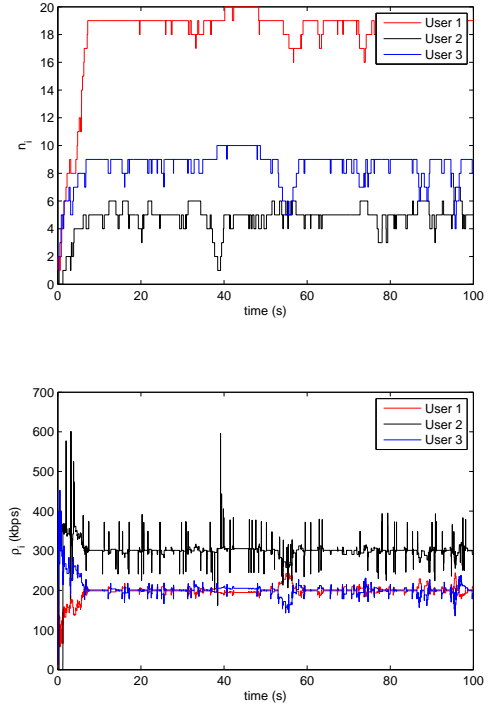


Fig. 2. Admission control in a linear network providing max-min fairness.

**Scenario 1.** In our first example we set the first link capacity to  $C_1 = 500\text{kbps}$ , the second link to  $C_2 = 400\text{kbps}$ , and the evolution of the system is simulated during  $100\text{s}$ . The arrival intensities of each user are  $\lambda_i = 2$  connections/second, and the mean workloads are  $1/\mu_i = 500\text{kbit} \approx 60\text{Kbytes}$ , making each link in the network congested.

We approximate max-min fairness by choosing user utilities in the  $\alpha$ -family with  $\alpha = 5$ . For the underlying congestion controllers, we choose  $\alpha = 2$  to emulate TCP-Reno behavior, with constants as in [11] inversely proportional to RTT. The long route has twice the RTT of the shorter routes.

The results in Figure 2 show that the user in the long route is allowed a greater number of flows, to compensate its RTT, and that aggregate loads of users 1 and 3 are almost equalized in  $\rho_i = 200\text{kbps}$  which is the fair share under max-min fairness. User 2 obtains the remaining bandwidth of link 1, that is  $300\text{kbps}$ . As its link is less congested, the user sees a lower congestion price, and therefore its connections run faster, which explains why admission control is letting less connections from this user enter the network.

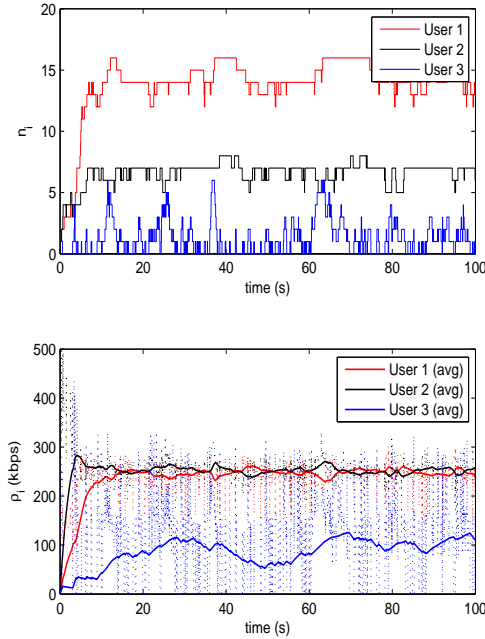


Fig. 3. Protection of non-greedy users in a linear network.

**Scenario 2.** We consider now the same linear network, but in a situation where there are heavy and non heavy users. By this we mean that some users demand an average load  $\lambda_i/\mu_i$  which is below the fair share they would obtain by solving Problem 1.

In our simulation, we achieve this by lowering the load of user 3 to  $\lambda_3 = 2$  connections/second, with mean workload  $1/\mu_3 = 50kbit$ . This amounts to an average load for user 3 of  $100kbps$ . The other users keep the same parameters as before.

In this situation, we would like User 3 to be *protected* by admission control: its connections should get through the second link because any congestion present in this link is due to User 1. This is in fact what happens in the simulation, as shown in Figure 3. Note that since user 3 is not congesting the link, its number of flows frequently hits zero, as does its aggregate rate. To visualize more clearly the utilization, we smoothed the aggregate rate obtained by each user.

As we can see, user 3 obtains approximately its average load  $100kbps$ , which is under the fair share the network will allocate in case of high demands. The remaining bandwidth is shared in a max-min way by greedy users, that is, they become bottlenecked in link 1 with  $\rho_1 \approx \rho_2 = 250kbps$ .

**Summary:** These simulations show that the proposed admission control can reestablish fairness at the user-level, by working with aggregate loads. It prevents greedy users from obtaining more than their fair share in the network; and if some users do not have enough traffic to reach their fair share, they obtain what they demand and the remaining capacity is allocated to the rest. As a final remark, the results are similar if we choose workload sizes from a heavy-tailed distribution, or if we use a different utility function for aggregate loads.

## V. CONCLUSIONS AND FUTURE WORK

The mathematics of Internet resource allocation made large strides with the introduction of microeconomic models [6], suggesting there could be a meeting point of the engineering side of congestion control and the economic side of user incentives. Still, the two areas have remained largely separate.

This paper contributes to bridging the gap by modelling user behavior in terms of the aggregate rate of all its connections, and controlling this aggregate through a combination of rate control of each connection, and admission control of the number of connections. Based on fluid-flow models, we developed a control law for the number of flows, which combined with standard TCP congestion control, has as equilibrium point the solution to an optimal user-level resource allocation. We have shown some theoretical results on the stability of this point.

The continuous control law guides us to propose an admission control policy for discrete flows, which leads to a Markov chain model for the number of flows in the network. In contrast to [1], [4], this procedure guarantees stability of the network for any external demand. Simulation evidence indicates this law approaches the optimal conditions of the fluid model; we are working on establishing this fact analytically.

An interesting extension we are currently pursuing is the combination of the flow dynamics with multipath routing. A multipath ingress router can not only decide whether to admit a connection or not, but also through which of the active routes to send it, based on the cheapest congestion price. Preliminary evidence shows this reaches the optimal rates of the multipath network utility maximization problem.

## REFERENCES

- [1] T. Bonald and L. Massoulié. Impact of fairness on internet performance. In *ACM SIGMETRICS/Performance*, pages 82–91, 2001.
- [2] B. Briscoe. Flow rate fairness: Dismantling a religion. *Computer Communication Review*, 37(2):63–74, 2006.
- [3] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. In *Proc. of the IEEE*, vol. 95, pages 255–312, Jan 2007.
- [4] G. de Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting services with rate control - could the internet be unstable? In *IEEE/Infocom*, pages 802–810, 1999.
- [5] S. B. Fredj, T. Bonald, A. Proutière, G. Régnié, and J. Roberts. Statistical bandwidth sharing: a study of congestion at flow level. In *ACM/SIGCOMM*, August 2001.
- [6] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 39:237–252, 1998.
- [7] H. K. Khalil. *Nonlinear Systems*. Prentice-Hall, Upper Saddle River, NJ, 1996.
- [8] X. Lin and N. B. Shroff. The impact of imperfect scheduling on cross-layer congestion control in wireless networks. *IEEE/ACM Transactions on Networking*, 14(2):302–315, 2007.
- [9] S. H. Low, F. Paganini, and J. C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 22(1):28–43, Feb. 2002.
- [10] K. Ma, R. Mazumdar, and J. Luo. On the performance of primal/dual schemes for congestion control in networks with dynamic flows. In *IEEE/Infocom*, 2008.
- [11] M. Mathis, J. Semke, J. Mahdavi, and T. Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *ACM Computer Communication Review*, 27(3):525–536, July 1997.
- [12] J. Mo and J. Walrand. Fair end-to-end window based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, Oct. 2000.
- [13] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston, MA, 2004.
- [14] J. T. Wen and M. Arcak. A unifying passivity framework for network flow control. In *IEEE/Infocom*, August 2003.