

Network resource allocation for users with multiple connections: fairness and stability

Andrés Ferragut and Fernando Paganini, *Senior Member, IEEE*

Abstract—This paper studies network resource allocation between users that manage multiple connections, possibly through different routes, where each connection is subject to congestion control. We formulate a user-centric Network Utility Maximization problem that takes into account the aggregate rate a user obtains from all connections, and propose decentralized means to achieve this fairness objective. In a first proposal, cooperative users control their number of active connections based on congestion prices from the transport layer, to emulate a suitable primal-dual dynamics in the aggregate rate; we show this control achieves asymptotic convergence to the optimal user-centric allocation. For the case of non-cooperative users, we show that network stability and user-centric fairness can be enforced by a utility-based admission control implemented at the network edge. We also study stability and fairness issues when routing of incoming connections is enabled at the edge router. We obtain in this case a characterization of the stability region of loads that can be served with routing alone, and a generalization of our admission control policy to ensure user-centric fairness when the stability condition is not met. The proposed algorithms are implemented at the packet level in ns2 and demonstrated through simulation.

I. INTRODUCTION

THE issue of fairness in resource allocation is fundamental to any shared infrastructure; as such it appears naturally in telecommunication networks. An important question in the network case is at which level of granularity or protocol layer should fairness be imposed. The main trend in networking research in recent times has been to seek fairness in the transport layer, between the allocated rates of end-to-end flows (or connections) traversing a network. Following the seminal work of Kelly et al. [17], this problem can be framed in terms of Network Utility Maximization (NUM), which captures various fairness notions between flows, including simplified yet powerful models of deployed TCP congestion control mechanisms, see [35]. The success of this methodology has projected NUM also into lower layers (routing, medium access, etc.), as a unifying technique to encompass multiple control mechanisms under a common fairness goal, see [6].

From the standpoint of network users, however, is the resulting fairness notion adequate? On the contrary, it appears that a *higher* layer aspect interferes: users can open an arbitrary number of connections across the network, skewing the overall rate allocation. In fact, aggressive applications often use this technique to vie for a larger share of the bandwidth “pie”, but even non-strategic users who happen to overload a common resource will be rewarded by a higher allocation. Therefore,

as argued in [3], we must go beyond flow-rate fairness for a more relevant view of network resource allocation.

In this paper, the object of fair allocation is a set of *users*; by definition, each user owns a set of connections, possibly through different routes. Our goal of network efficiency and fairness is the rate allocation between users, in a manner that optimizes a *user-centric* NUM problem. To achieve this objective, we propose to actively control the number of flows per user, assuming the underlying per-flow allocation is unchanged from the aforementioned standard models of congestion control. We now outline our contributions, other related work is summarized in Section II.

Our first result, presented in Section III, is related to the motivation of users to increase the number of active flows. We show, under fairly general assumptions on the network topology, that the aggregate rate a user obtains in a certain route increases with the number of connections in this route, when the competing connections are fixed; thus users’ selfish incentives are aligned with increasing connection numbers beyond limit, a mutually destructive outcome.

Achieving user-centric fairness therefore requires controlling connection numbers; in Section IV we analyze whether this objective is achievable in a decentralized fashion, assuming temporarily that users are cooperative. Since connections may use different routes, the required dynamics of aggregate rates are of the form of multi-path congestion control, which is well known to suffer from oscillations. We propose for this purpose a new variant of primal-dual congestion control which is shown to be globally asymptotically stable, and is well suited for implementation through a connection-level dynamics, using available congestion feedback from the network.

Since user cooperation cannot be counted on, in Section V we propose a decentralized admission control rule, based on user utilities and thus tailored to our proposed user-centric fairness. We analyze the performance of this control under a traffic model of random connection arrival/departures, through a fluid limit argument. The mechanism is shown to protect the network from greedy users, imposing in situations of overload the desired notion of fairness.

In Section VI we turn our attention to the related problem of connection-level routing: users bring end-to-end jobs to transfer, with routes chosen by the network. While each individual connection remains single-path, users may now profit from several routes. We characterize the stability region of this problem, and give conditions under which it is attainable by a simple congestion-based routing policy. We also show how to combine admission control and routing to provide stability and fairness when the loads exceed the natural stability region.

Finally, we provide in Section VII packet-level simulations to test the proposed algorithms in practice. The algorithms are implemented in `ns2`, either at the end-hosts for the cooperative control case, or at the edge router in the admission control/routing cases. Our simulations validate the accuracy of our model predictions, in particular exhibiting the desired fairness. Conclusions are presented in Section VIII, and an Appendix contains some of the proofs. Partial versions of these results were presented in [10], [11]; for more extensive details we refer to the thesis [9].

II. RELATED WORK

Our work touches on several topics that have been studied in other references; these are now overviewed.

The impact of parallel TCP connections on aggregate throughput is analyzed in [13], experimentally and invoking the TCP rate formulas of [30]. In [40] these formulas are used for an analysis of strategic user incentives in a single bottleneck network. Our analysis, based on the NUM framework, enables us to generalize the conclusions to arbitrary network topologies, as well as different notions of flow-rate fairness.

Multi-path congestion control involves endowing each end-to-end connection with multiple paths over which to send traffic, with the capability of controlling each path rate. This has been analyzed from a theoretical perspective in the NUM setting already in [17], for so-called primal algorithms that solve a barrier approximation to NUM; [14], [18] later analyzed the delay stability of this solution. For the exact NUM problem, the difficulty that appears is the lack of strict concavity of the objective function, which leads to oscillations in gradient-type methods. In this respect, the pure dual algorithm considered in [36] yields a discontinuous dynamics that chatters around the equilibrium value, converging only in a mean sense. In [37] this is addressed by replacing the objective function by a strictly concave approximation, thus leading to a stable approximate algorithm. Another strategy to obtain strict concavity is the so-called proximal optimization method, which was applied to multi-path TCP in [25], leading to discrete time algorithms that converge under suitable step size conditions. Non-strict concavity also compromises stability of primal-dual control laws (see [8]); in this regard, our proposal of Section IV provides a new, globally convergent primal-dual law that could be applied to the multi-path TCP problem. From a practical perspective, there is an ongoing discussion in the IETF on multi-path TCP implementations, see e.g. [23], [33], [39]. In contrast to these transport layer implementations, our main motivation here is to use the analysis as a basis for controlling the *number* of (individually single path) connections to achieve efficiency and fairness in the aggregate rates.

The use of connection-level control to modify the resource allocation provided by the network was proposed in [4], [5], in the context of wireless networks. Motivated by the high loss rate in these environments, which tampers with adequate congestion feedback, the authors propose an Inverse-Increase Multiplicative-Decrease algorithm to adjust the number of connections, an application layer strategy that imposes a certain resource allocation on the problem, overcoming the

lossy wireless channel. Our results of Section IV rely on the same type of control but take the strategy further, to impose an arbitrary desired fairness model on the aggregate rates of a set of users over possibly multiple paths. This proposal is philosophically aligned with the suggestion of [17] that user-specific utilities can be reconciled with congestion control protocols by adjusting a weight parameter in the latter. However, adjusting the number of connections is more amenable to implementation at the application layer, without changing the current transport layer. A recent reference on the latter strategy is [38]. Our approach has similarities to the “coordinated congestion control” studied in [21], but there are differences in the optimization objective sought and the connection dynamics considered.

Another way to take connection dynamics into account is through a queuing model for network flows, modeled by stochastic processes or their fluid limits, for which TCP resource allocation is a service discipline. In this line, [1], [7] showed that the natural stability condition (all average link loads less than their capacity), is indeed sufficient for stability in the memoryless case. This analysis has been extended in several ways in [24], [26], [32] to more general hypotheses, particularly in the job sizes. In [14], [20], the corresponding conditions were given for operation under multi-path TCP; [20] also shows that an “uncoordinated” control of single-path connections may not in general be able to stabilize the complete region. In our work of Section VI we also employ single-path connections, but we add congestion-based routing in a way that allows us to cover the full stability region. Other related work on connection routing is [15], where optimal routing policies are obtained under the assumption that the network provides a so-called balanced fair allocation; this however does not apply to typical congestion control protocols.

Note, finally, that such stochastic stability results are of an open-loop nature: either the loads are stabilized and users are satisfied, or the network is unstable, and this is independent of the congestion control applied. Some authors [16], [27] have argued from here that admission control of connections is required. While any reasonable admission control may overcome such instability by discarding excess connections, the distinguishing feature of our utility-based admission control of Section V is that a desired *fairness* between users is imposed in such situations of overload.

III. FLOW-LEVEL FAIRNESS LIMITATIONS

We consider a network composed of links, indexed by l , with capacity c_l , and a set of paths or routes, indexed by r . End-to-end connections (flows) travel through a single path, specified by the routing matrix R ($R_{lr} = 1$ if route r contains link l , and 0 otherwise). x_r denotes the rate of a single connection along route r . Let n_r denote the number of such connections, with $\varphi_r = n_r x_r$ denoting the aggregate rate. The rate through link l can be expressed as $y_l = \sum_r R_{lr} \varphi_r = \sum_r R_{lr} n_r x_r$.

Connections present in the network regulate their rate through some congestion control mechanism, which we model (c.f. [35]) as seeking the solution of the following convex optimization:

Problem 1 (Congestion control): For fixed $\{n_r\}$, $n_r > 0$,

$$\max_{\varphi_r} \sum_r n_r U_{TCP_r} \left(\frac{\varphi_r}{n_r} \right),$$

subject to the capacity constraints $y_l \leq c_l \forall l$.

The above optimization provides a notion of “flow-rate fairness”, where U_{TCP_r} reflects the choice of the congestion controller,¹ and this utility is assigned to the individual connection rate $x_r = \varphi_r/n_r$. These utilities are assumed increasing and strictly concave; we focus here on the usual α -fair family [29], which satisfies $U'_{TCP_r}(x_r) = w_r x_r^{-\alpha}$, and encompasses many commonly used fairness models.

Decentralized methods to solve Problem 1 involve the use of duality. Let p_l denote the Lagrange multipliers (prices) associated with each link constraint, and q_r denote the aggregate route prices

$$q_r = \sum_l R_{lr} p_l. \quad (1)$$

The Karush-Kuhn-Tucker (KKT) conditions for Problem 1 include $U'_{TCP_r}(\varphi_r/n_r) = q_r$, equivalent to the *demand curve*:

$$x_r = \frac{\varphi_r}{n_r} = f_{TCP_r}(q_r), \quad (2)$$

with $f_{TCP_r} = [U'_{TCP_r}]^{-1}$. In particular, for α -fair utilities $f_{TCP_r}(q_r) = w_r q_r^{-1/\alpha}$.

Therefore, congestion control algorithms behave as decentralized ways to solve Problem 1, where U_{TCP_r} reflects *protocol behavior*. This in turn defines a mapping $\Phi : n \mapsto \varphi$ where, given the number of connections $n = (n_r)$ in each route, the resource allocation $\varphi = (\varphi_r)$ is calculated as the solution of the Congestion Control Problem 1. From a user perspective, for a given number of connections in each route, the allocated resources are determined by this flow-rate fairness. However, a user vying for more resources may challenge this by opening more connections. We have the following result, proved in the Appendix:

Theorem 1: Assume that R has full row rank. Then the map $\varphi = \Phi(n)$ above is such that:

$$\frac{\partial \varphi_r}{\partial n_r} \geq 0 \quad \forall r, n.$$

This result implies that greedy users have incentives to increase the number of connections to bias the resource allocation over any network topology.² Moreover this holds independently of the utility used by the congestion control layer, i.e. the underlying algorithm. Formalizing this further, assume each user has an increasing and concave utility $U_i(\varphi^i)$ modeling its valuation of the total rate φ^i obtained from the network, and each strategically chooses the number of connections. In this connection game, Theorem 1 implies that a dominant strategy is to increase the number of connections along any route. If a subset of greedy users behave in this way, the number of ongoing connections will grow without

bounds, an undesirable scenario. This formalizes and exhibits the limitations of flow rate fairness mentioned in [3].

A non strategic way of taking users into account is through stochastic models for demand. Here connections arrive on route r as a Poisson process of intensity λ_r , with each connection bringing a random amount of workload with mean $1/\mu_r$. For each network state (n_r) , rates are assigned according to Problem 1. This model was first analyzed in [1], [7], where under the hypothesis of exponentially distributed workloads, the stochastic process n_r is stable provided:

$$\sum_r R_{lr} \rho_r < c_l \quad \forall l \quad (3)$$

where $\rho_r = \lambda_r/\mu_r$ is the average load on route r . This stability condition has also been extended in different ways in [25], [26], [32], in particular to general workload distributions.

The stochastic stability of this system is therefore characterized. However, congestion control plays no role in enforcing stability: if (3) is not satisfied, the number of ongoing connections will grow without bounds, up to a point where user impatience comes into play and connections are dropped. Some authors [16], [27] argued that the above situation requires admission control of connections. While simple admission control rules may overcome instability, the remaining question is how to carry it out in a way that fairness *between users* is taken into account. We now investigate further this notion of fairness.

IV. USER-CENTRIC FAIRNESS OVER MULTIPLE PATHS

Assume that there is a set of *users*, indexed by i , which open connections in the network. Each user therefore has a set of routes r and receives an aggregate rate of service $\varphi^i = \sum_{r \in i} \varphi_r$. Let U_i be an increasing and concave utility function that models user preferences instead of protocol behavior. The associated *user-centric* notion of fairness can be expressed through the following NUM problem:

Problem 2 (User Welfare):

$$\max_{\varphi_r} \sum_i U_i(\varphi^i)$$

subject to link capacity constraints $y_l = \sum_r R_{lr} \varphi_r \leq c_l \forall l$. Here, the sum in the constraints is done over all the network routes. Each route is associated with a single user, and if several users open connections along the same path, we duplicate the index r accordingly. Note also that the above framework is very general, with a user defined as a set of routes. This can model users downloading data from several locations, multiple parallel paths, the single-path case, etc.

A first step in our analysis will be to assume that users cooperate by controlling the aggregate rate on each route: we will construct a dynamics for the φ_r that globally drive the system to the desired optimum, and then analyze how to implement it through connection level control. Consider the Lagrangian of Problem 2:

$$\begin{aligned} L(\varphi, p) &= \sum_i U_i(\varphi^i) - \sum_l p_l (y_l - c_l) \\ &= \sum_i U_i \left(\sum_{r \in i} \varphi_r \right) - \sum_r q_r \varphi_r + \sum_l p_l c_l. \end{aligned} \quad (4)$$

¹To simplify the notation, we use U_{TCP} for the utility of the lower layer congestion controller, of which TCP is a particular case.

²The hypothesis on R is typically a non-issue since there are more routes than links in the network.

The KKT conditions that characterize the saddle point of L on (4) imply that for each r ,

$$\begin{aligned} &\text{either } U_i' \left(\sum_{r \in i} \varphi_r \right) = q_r^*, \\ &\text{or } U_i' \left(\sum_{r \in i} \varphi_r \right) < q_r^* \text{ and } \varphi_r^* = 0. \end{aligned} \quad (5)$$

In particular, (5) implies that $U_i'(\varphi^{i,*}) = q_i^* := \min_{r \in i} q_r^*$, i.e. user i only sends traffic through minimum price paths. While $\varphi^{i,*}$ is determined by the above, the optimal rates φ_r^* need not be unique. Problem 2 coincides with the multi-path congestion control problem considered in [36], so a number of distributed approaches are available to drive φ_r to the optimum. However, some difficulties appear due to the lack of strict concavity in the objective, which often leads to oscillatory behavior. Our proposal is to use a variant of the primal-dual dynamics with an additional damping term to obtain convergence. Consider the following control law:

$$\dot{\varphi}_r = k_r (U_i'(\varphi^i) - q_r - \nu \dot{q}_r)_{\varphi_r}^+, \quad (6a)$$

$$\dot{p}_l = \gamma_l (y_l - c_l)_{p_l}^+, \quad (6b)$$

where $y = R\varphi$, $q = R^T p$ as before, and the gains $k_r, \gamma_l, \nu > 0$. $(\cdot)_v^+$ is the positive projection, which verifies $(u)_v^+ = 0$ whenever $v = 0$ and $u < 0$, otherwise $(u)_v^+ = u$.

Algorithm (6) is a modified primal-dual algorithm in which end users adjust their rates according to a *predicted* route price $q_r + \nu \dot{q}_r$, thus anticipating possible changes. This idea first appeared in [31] in the context of combined multi-path congestion control and routing. Note that this damping term does not affect the equilibrium, and due to the use of the projection, the equilibrium of (6) verifies the KKT conditions (5). We have the following:

Theorem 2: Under the control laws given in equation (6), all trajectories converge to a solution of Problem 2.

The proof is based on the following Lyapunov function:

$$V(\varphi, p) = \sum_r \frac{(\varphi_r - \varphi_r^*)^2}{2k_r} + \sum_l \frac{(p_l - p_l^*)^2}{2\gamma_l} + \nu (c_l - y_l^*) p_l, \quad (7)$$

where (φ^*, p^*) is an equilibrium and $y^* = R\varphi^*$. Note that $V \geq 0$ for every $\varphi, p \geq 0$, in particular the last term is non-negative due to the Problem constraints $y_l^* \leq c_l$ and $p_l \geq 0$. Also, this term vanishes in any equilibrium due to (6b), which in turn imposes the complementary slackness condition. The full derivation is presented in the Appendix.

Theorem 2 shows that the dynamics (6) become a good alternative for multi-path congestion control. This algorithm is decentralized, since it only assumes that user i can control the rate on its own routes, using only the total route price and its derivative. However, instead of changing congestion control procedures, we would like to derive connection-level controllers that use the number of ongoing connections to drive the system to equilibrium. This way, individual connections relay on current transport layer protocols, which hide the network complexity. The application layer then controls φ_r only indirectly through the number of connections. We now address this issue.

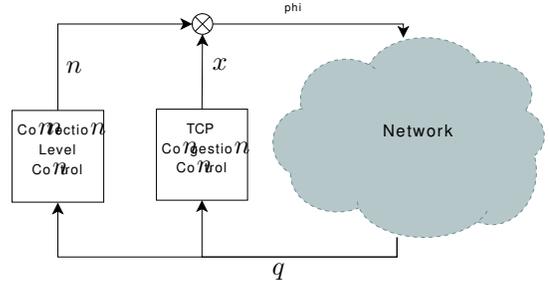


Fig. 1. Block diagram of the proposed connection level control

A. Connection-level control

We propose now a connection level dynamics for n_r , the number of connections on each route. In order to achieve this, we shall consider a way to control n_r such that $\varphi_r = n_r x_r$ follows equation (6). Note that in this case $\dot{\varphi}_r = \dot{n}_r x_r + n_r \dot{x}_r$. Consider the following control law:

$$\dot{n}_r = n_r \left[k (U_i'(\varphi^i) - q_r - \nu \dot{q}_r) - \frac{\dot{x}_r}{x_r} \right]. \quad (8)$$

With this choice, it is easy to check that $\dot{\varphi}_r = k \varphi_r [U_i'(\varphi^i) - q_r - \nu \dot{q}_r]$, which is similar to (6) but with a state-dependent gain $k \varphi_r$.

The problem under consideration is best explained through Fig. 1: on the right, we represent the network by an entity that receives aggregate rates φ_r , and returns congestion prices q_r per route. These are used by congestion control to generate the rate x_r per connection; thus the inner loop represents TCP congestion control, for fixed n_r . What we wish to design here is the outer loop (which operates at a slower time-scale), controlling the n_r such that the overall dynamics of φ_r achieves the desired user-centric fairness.

For further clarity, and to facilitate implementation, it is convenient to rewrite the dynamics of n_r in terms of the congestion price, eliminating the variable x_r . Assume that congestion control can be modeled using utilities from the α -fair family, we have $U_{TCP_r}^l(x_r) = w_r x_r^{-\alpha}$ or equivalently $x_r = f_{TCP_r}(q_r) = (q_r/w_r)^{-1/\alpha}$. The last term can be rewritten as

$$\frac{\dot{x}_r}{x_r} = -\frac{1}{\alpha} \frac{q_r^{-1/\alpha-1} \dot{q}_r}{q_r^{-1/\alpha}} = -\frac{\dot{q}_r}{\alpha q_r}$$

and therefore the dynamics of n_r becomes

$$\dot{n}_r = n_r \left[k (U_i'(\varphi^i) - q_r - \nu \dot{q}_r) + \frac{\dot{q}_r}{\alpha q_r} \right]. \quad (9)$$

We have the following:

Proposition 1: The connection-level dynamics (9) globally stabilizes the equilibrium of Problem 2.

The proof is based on a Lyapunov function similar to (7), with a minor modification to account for the state-dependent gain $k \varphi_r$ (see Remark 3 in the Appendix). Observe also that in equation (9), the predictive terms in \dot{q}_r play opposing roles. This suggests considering the simpler control law

$$\dot{n}_r = k n_r (U_i'(\varphi^i) - q_r); \quad (10)$$

in fact when translated to φ_r , this yields dynamics very similar to (6), with derivative action in the control of $\dot{\varphi}_r$, except that the damping parameter $\nu_r = (\alpha k q_r)^{-1}$ is time varying and route dependent, which is not compatible with our earlier stability argument. Extending the proof to this case remains open at this time. For the single-path case, we proved in [10] that the above dynamics are locally asymptotically stable.

Also, the control law (9), while having guaranteed global properties, uses the derivative action term which can be hard to implement in practice due to noisy measurements of the price. The simpler law (10) is more amenable to implementation, and extensive (both fluid and packet level) simulations show that it is well behaved [9]. In Section VII we shall explore a packet-level implementation of the latter mechanism.

V. UTILITY BASED ADMISSION CONTROL

The model in Section IV is applicable to the case where users cooperate by opening or closing connections based on an appropriate feedback from the network. Since selfish incentives of users do not encourage this behavior, we cannot generally count on this cooperation. In such cases, the network must take the role of controlling connection numbers, for which the simplest means is *admission control*. This approach was advocated in [27], where a stochastic model of connection arrivals and departures is discussed, and admission control is used to ensure the stochastic stability of the system when the average load is larger than the link capacity; this is done without addressing fairness in the resulting resource allocation. We now would like to derive a decentralized admission control rule, that can be enforced at the network edge, and such that in case of overload resources are allocated according to the User Welfare Problem 2.

From our analysis in Section IV, we see that in order to achieve fairness, each user must increase its number of connections whenever $U'_i(\varphi^i) > q_r$, i.e. the user marginal utility is greater than the current congestion price. If on the other hand the inequality reverses, the number of connections must be decreased. Consider the following admission control rule for incoming (new) connections:

$$\begin{aligned} \text{If } U'_i(\varphi^i) > q_r &\rightarrow \text{admit connections on route } r. \\ \text{If } U'_i(\varphi^i) \leq q_r &\rightarrow \text{drop connections on route } r. \end{aligned} \quad (11)$$

where $\varphi^i = \sum_{r \in i} \varphi_r$, as before. We call this rule *Utility Based Admission Control*. Equation (11) imposes a limit on the number of connections a given user is allowed, and therefore a strategic user will not get a larger share of bandwidth simply by opening more connections, as in Theorem 1, since eventually the admission condition will not be met. In a scenario where all users are pushing the limits, the network will operate in the region $U'_i(\varphi^i) \approx q_r$ for all r, i , which are the KKT conditions of Problem 2.

We now formalize these arguments using a stochastic model for the system. We discuss first the single-path case, and postpone the discussion of the multi-path case to Section V-C.

A. Admission control in the single path case

In the single path case, each user i is associated with a single route r , and thus we can write U_r for the user utility

function instead of U_i . In this case, the rule (11) reduces to:

$$\begin{aligned} \text{if } U'_r(\varphi_r) > q_r &\rightarrow \text{admit connection,} \\ \text{if } U'_r(\varphi_r) \leq q_r &\rightarrow \text{drop connection.} \end{aligned} \quad (12)$$

Assume each user on route r opens connections, which arrive as a Poisson process of intensity λ_r , and bring an exponentially distributed workload of mean $1/\mu_r$. Connection arrival and job sizes are independent and also independent between users. Assuming a time scale separation, i.e. that congestion control operates faster than the connection level process, the rate at which a connection is served is $x_r(n) = \varphi_r(n)/n_r$, determined by the solution of Problem 1. Also, the aggregate rate on route r is $\varphi_r(n)$ and $q_r(n)$ is the route price. This model was introduced by [1], [7]. When the admission control rule (12) is added, the process $n(t)$ is a continuous time Markov chain with the following transition rates:

$$\begin{aligned} n &\mapsto n + e_r && \text{with rate } \lambda_r \mathbf{1}_{\{U'_r(\varphi_r) > q_r\}}, \\ n &\mapsto n - e_r && \text{with rate } \mu_r \varphi_r, \end{aligned} \quad (13)$$

where e_r is the vector with a 1 in coordinate r and 0 elsewhere, and $\mathbf{1}_A$ is the indicator function.

Without the admission condition, [1], [7] prove that the Markov chain is stable (positive recurrent) if the loads $\rho_r = \lambda_r/\mu_r$ satisfy the natural condition

$$\sum_r R_{lr} \rho_r < c_l \quad \text{for each } l. \quad (14)$$

On the other hand, admission control should stabilize the system in any situation. This is indeed the case for rule (12).

Proposition 2: The Markov chain given by (13) is stable.

Proof sketch: The proof relies on constructing a suitable ‘‘box’’ set $S = \{n : \|n\|_\infty \leq n_0\}$ with n_0 large enough such that, if $n_r = n_0$, then the admission condition is violated. Therefore, the process starting at an empty network cannot leave S , and since the Markov chain is irreducible, it will converge to an equilibrium distribution on a subset of the finite set S . For details see [9]. ■

B. Fluid limit analysis

Now that stability is assured, we proceed to analyze the fairness of the admission control policy. We will do so by deriving a suitable fluid model for the system (13). The model is based in a *large network asymptotic*. The main idea is to scale the network size appropriately, by enlarging the capacity of the links and the arrival rate of flows, such that a law of large numbers scaling occurs. An important remark is that, for the scaling to work appropriately, we also have to scale the user preferences with the size of the network.

More formally, we take a scaling parameter $L > 0$ and consider a network with link capacities scaled by L , i.e. $c_l^L = Lc_l$. We also assign each user a utility $U_r^L(\varphi) = LU_r(\varphi/L)$. Note that with this choice, the utility functions verify the following scaling property:

$$(U_r^L)'(L\varphi) = U'_r(\varphi).$$

That is, the user marginal utility for obtaining L times bandwidth in the scaled network is the same that the marginal utility for the original amount in the original network.

We denote by $\varphi^L(n)$ and $q^L(n)$ the rate allocation and route prices in the scaled network, and as before $\varphi(n)$ and $q(n)$ denote the original values, i.e. $L = 1$. The following relationships are direct from the KKT conditions:

Lemma 1: For any $L > 0$, the resource allocation and route prices satisfy:

$$\varphi_r^L(Ln) = L\varphi_r(n), \quad q_r^L(Ln) = q_r(n) \quad \forall r.$$

Using the above relationships, we now derive the fluid model of the system. To avoid technicalities, we shall replace the indicator function in (13) by a smooth approximation f_ϵ , such that $f_\epsilon(x) = 1$ if $x > \epsilon$, and $f_\epsilon(x) = 0$ if $x < -\epsilon$. The original model can therefore be approximated by:

$$\begin{aligned} n &\mapsto n + e_r && \text{with rate } \lambda_r f_\epsilon(U_r'(\varphi_r) - q_r), \\ n &\mapsto n - e_r && \text{with rate } \mu_r \varphi_r. \end{aligned} \quad (15)$$

Note that, as $\epsilon \rightarrow 0$ the above model approximates (13). We have the following result, proved in the Appendix:

Theorem 3: Consider a sequence of processes $n^L(t)$ governed by equations (15) with $\lambda_r^L = L\lambda_r$, $c_l^L = Lc_l$, $\mu_r^L = \mu_r$, and utility functions that satisfy

$$U_r^L(\varphi) = LU_r\left(\frac{\varphi}{L}\right),$$

with $L > 0$ a scaling parameter. Consider also a sequence of initial conditions $n^L(0)$ such that $n^L(0)/L$ has a limit $\bar{n}(0) > 0$. Then, the sequence of scaled processes:

$$\bar{n}^L(t) = \frac{n^L(t)}{L}$$

is tight and any weak limit point of the sequence converges as $L \rightarrow \infty$ to a solution of the following differential equation:

$$\dot{n}_r = \lambda_r f_\epsilon(U_r'(\varphi_r) - q_r) - \mu_r \varphi_r, \quad \text{for } n_r > 0, \quad (16)$$

where $\varphi(n)$ and $q(n)$ are the allocation maps for $L = 1$.

Remark 1: We have constrained the dynamics (16) to the region $n_r > 0$; a complete model would require describing what happens if the trajectory reaches the boundary. In this regard, we note that, if the user average loads $\rho_r = \lambda_r/\mu_r$ satisfy the natural stability condition (14), it is shown in [1] that the trajectory $\{n_r\}$ reaches zero in finite time. When $n_r = 0$ for some r , the allocation φ_r drops to zero, but then the arrivals term in (16) would move the state back into positive values, and then back to zero as service rate appears. Thus the state will remain ‘‘chattering’’ around zero, and should receive an average service rate ρ_r , a behavior that is difficult to express precisely in ordinary differential equation terms; see [19] for fluid models that take this aspect into account.

Here, we are mainly interested in the case of overload, where at least one link capacity is exceeded, and admission control must apply to some of the users. There might be certain users that are completely isolated from this overload, using only non-congested routes; these would be stabilized as discussed above. Therefore, without loss of generality we assume henceforth that all participating users share at least one of the overloaded links; in that case the equilibrium of (16) will have nonzero n_r for all r 's; we can thus avoid boundary effects when analyzing the local dynamics around equilibrium.

The equilibrium condition for (16) is:

$$\varphi_r^* = \frac{\lambda_r}{\mu_r} f_\epsilon(U_r'(\varphi_r^*) - q_r^*) = \rho_r f_\epsilon(U_r'(\varphi_r^*) - q_r^*).$$

Since $\varphi_r^* > 0$ and f_ϵ is bounded above by 1, the only possibilities are:

$$\begin{aligned} \varphi_r^* &< \rho_r && \text{and } |U_r'(\varphi_r) - q_r| < \epsilon && \text{or} \\ \varphi_r^* &= \rho_r && \text{and } U_r'(\varphi_r) > q_r + \epsilon. \end{aligned}$$

As $\epsilon \rightarrow 0$ the above translate to:

$$\varphi_r^* < \rho_r \quad \text{and} \quad U_r'(\varphi_r^*) = q_r^* \quad \text{or} \quad (17a)$$

$$\varphi_r^* = \rho_r \quad \text{and} \quad U_r'(\varphi_r^*) > q_r^*. \quad (17b)$$

The interpretation of the above conditions is the following: either the equilibrium allocation for user r is less than its demand, and the system is on the boundary of the admission condition, or the user is allocated its full average demand and admission control is not applied.

We would like to relate this to the User Welfare Problem 2 defined before. Considered the following:

Problem 3 (Saturated User Welfare):

$$\max \sum_r U_r(\varphi_r),$$

subject to $R\varphi \leq c$ and $\varphi_r \leq \rho_r \quad \forall r$.

Problem 3 has the following interpretation: allocate resources to the different users according to their utility functions, but do not allocate a user more than its average demand. It amounts to saturating the users to a maximum possible demand, given by the value ρ_r .

We have the following Proposition, whose proof is direct of the KKT conditions and equations (17):

Proposition 3: As $\epsilon \rightarrow 0$, the equilibrium points of (16) converge to the optimum of Problem 3.

Therefore, the equilibrium allocation under admission control in an overloaded network is a solution of Problem 3. Note that if traffic demands are very large ($\rho_r \rightarrow \infty$), Problem 3 becomes the original User Welfare Problem 2, and admission control is imposing the desired notion of fairness. Moreover, if some users demand less than their fair share according to Problem 2, the resulting allocation *protects* them from the overload by allocating these users their mean demand, and sharing the rest according to the user utilities. In Section VII we shall give a numerical example of this behavior.

C. Admission control in the multi-path case

Consider now the situation where the user opens connections on several paths, and obtains utility from the aggregate. Assume that connection arrivals on each path are independent, following a Poisson process of intensity λ_r , and with exponentially distributed workloads of mean $1/\mu_r$. For example, this would be the case of users downloading data from different sources at the same time. The lower layers of the network allocate resources as in the single path situation, and each route has an average load $\rho_r = \lambda_r/\mu_r$. Assume that the network implements the admission control rule (11), controlling the *aggregate rate* each user perceives.

In this case, by a similar analysis, the dynamics become $\dot{n}_r = \lambda_r f_\epsilon(U_i^i(\varphi^i) - q_r) - \mu_r \varphi_r$, and in the overload case, the dynamics converges to the solution of:

Problem 4 (Saturated User Welfare (multi-path)):

$$\max \sum_i U_i(\varphi^i)$$

subject to $R\varphi \leq c$, $\varphi_r \leq \rho_r \forall r$.

This in turn implies that the admission control rule (11) applied to the aggregate load can be effectively used to drive the network to a fair allocation, even when users demands are transported over different routes.

VI. CONNECTION-LEVEL ROUTING

The analysis of the preceding section assumes that each user establishes connections through some set of predefined routes, possibly with multiple destinations. The user manages simultaneously several connections over these routes and derives a utility from the aggregate rate. Moreover, the user has an independent arrival rate λ_r for each route. We now focus on a slightly different situation: here, each user has a set of routes available to communicate with a given destination in the network. These routes are indifferent for the user, all of them serving the same purpose. Each user brings connections into the network, and at each connection arrival, the user or the edge router may decide over which route to send the data. This is a typical instance of the multi-path load balancing problem, but at the connection level timescale.

We consider an adaptation of the stochastic model for connections of [1], [7] to this problem. Assume that users have multiple routes available to serve their jobs. User i generates incoming connections as a Poisson process of intensity λ_i , and exponential file sizes with mean $1/\mu_i$. Thus $\rho^i = \lambda_i/\mu_i$ represents the *user* average load. Here we do not distinguish between the routes, since each user may be served by a set \mathcal{R}_i of possible routes. As for congestion control, we assume that the TCP layer can be described as in the Congestion Control Problem 1. The user or the network may now choose, at the start of the connection, to which route to send the data from the set \mathcal{R}_i , but each connection behaves independently after that, following a single specified path throughout its service. Nevertheless, by appropriately choosing the route, the load may be distributed across multiple paths.

We formalize a routing policy in the following way: given the current state of the network, characterized by the vector $n = (n_r)$ of ongoing connections per route, a routing policy is a selection $r \in \mathcal{R}_i$ for a new connection. We denote by $A_{ir} \subset \mathbb{N}^N$ the set of network states such that connections arriving from user i are routed through route $r \in \mathcal{R}_i$. If the same physical route is possible for many users, we duplicate its index r accordingly, and N is the total number of routes.

The only general requirement for the routing policy is that the sets A_{ir} are a partition of the space for each i , i.e.:

$$\sum_{r \in \mathcal{R}_i} \mathbf{1}_{A_{ir}} \equiv \mathbf{1} \quad \forall i. \quad (18)$$

In a fluid limit, the dynamics of the number of connections under the routing policy $\{A_{ir}\}$ is given by:

$$\dot{n}_r = (\lambda_i \mathbf{1}_{A_{ir}} - \mu_i \varphi_r)_{n_r}^+. \quad (19)$$

Here, $\varphi_r = \varphi_r(n)$ as before denotes the total rate assigned to the flows on route r depending on network state. The saturation $(\cdot)_{n_r}^+$ is needed in this case because some routes may not be used, and thus the number of flows must remain at 0.

Remark 2: We could have also considered more general routing policies, in which each routing decision is assigned a probability $p_{ir}(n)$ for each network state. The routing policy constraint (18) in that case will be the same. However, in the following we will only focus on deterministic routing policies.

Note also that the sets A_{ir} may be quite general. However, for practical implementation, it is necessary that the routing policy is decentralized, i.e. the decision of routing a flow of user i over route r should be based on local information. We defer this discussion for the moment and focus on necessary conditions for stability of the system.

A. A necessary condition for stability

Our goal is to characterize the stability region of a routing policy, with dynamics given by (19). More formally, we would like to know for which values of ρ^i the fluid model goes to zero in finite time. We recall that finite time convergence of the fluid model is related (c.f. [34, Chapter 9]) to the stochastic stability of the corresponding Markov chain models. We will first derive a necessary condition for stability, which generalizes the stability condition of [1], [7] to the routing case.

For this purpose, introduce for each user i the simplex of possible traffic fractions among available routes:

$$\Delta^i = \left\{ \alpha^i = (\alpha_r^i)_{r \in \mathcal{R}_i} : \alpha_r^i \geq 0, \sum_{r \in \mathcal{R}_i} \alpha_r^i = 1 \right\}.$$

The following is the main result of this section:

Theorem 4: A necessary condition for the existence of a policy $\{A_{ir}\}$ that stabilizes the dynamics (19) is that, for each user i , there exists a split $\alpha^i \in \Delta^i$ such that:

$$\sum_l R_{lr} \alpha_r^i \rho^i \leq c_l \quad \forall l. \quad (20)$$

Condition (20) is the non-strict version of single-path condition (14) for the split traffic loads $\rho_r = \alpha_r^i \rho^i$. Thus, equation (20) means that for a routing policy to exist, it is necessary that the network is “stabilizable”, in the sense that there is a partition of the user loads such that the underlying single path network is stable. Of course, if each user has only one possibility, then $\Delta^i = \{1\}$ and we recover the single path stability condition. The same condition (20) was obtained in [14] for stochastic stability in the case of multi-path TCP. In that case, however, the TCP layer must be modified to make full simultaneous use of the available routes. Here each route remains single-path, with standard congestion control, and the routing policy is used to achieve the same stability region, whenever possible.

Proof of Theorem 4: Consider the convex and compact subset of \mathbb{R}^L , with L the total number of links:

$$\mathcal{S} = \left\{ z \in \mathbb{R}^L : z_l = \sum_r R_{lr} \alpha_r^i \rho^i - c_l, \alpha^i \in \Delta^i \right\}$$

The set \mathcal{S} represents the excess of traffic in each link for each possible split. If (20) is not feasible for a given load vector ρ^i , then the set \mathcal{S} is disjoint with the closed negative orthant \mathbb{R}_-^L . By convexity, there exists a strictly separating hyperplane [2, Section 2.5], i.e. a fixed $\hat{p} \in \mathbb{R}^L$ with $\hat{p} \neq 0$ and such that:

$$\begin{aligned} \hat{p}^T z &\geq a + \epsilon \quad \forall z \in \mathcal{S}, \\ \hat{p}^T z &\leq a \quad \forall z \in \mathbb{R}_-^L. \end{aligned}$$

The second condition implies in particular that \hat{p} has non-negative entries, since if $\hat{p}_l < 0$, taking $z_l \rightarrow -\infty$ we have $\hat{p}_l z_l \rightarrow +\infty$ and therefore the inequality is violated for any a . Also, since $z = 0 \in \mathbb{R}_-^L$, we have that $a \geq 0$. Define now $\hat{q} = R^T \hat{p}$ and consider the following state function:

$$V(n) = \sum_i \frac{1}{\mu_i} \sum_{r \in \mathcal{R}_i} \hat{q}_r n_r. \quad (21)$$

Note that $\hat{q}_r \geq 0$ and $\hat{q} \neq 0$.

Differentiating V along the trajectories of (19) we get:

$$\begin{aligned} \dot{V} &= \sum_i \frac{1}{\mu_i} \sum_{r \in \mathcal{R}_i} \hat{q}_r \dot{n}_r = \sum_i \frac{1}{\mu_i} \sum_{r \in \mathcal{R}_i} \hat{q}_r [\lambda_i \mathbf{1}_{A_{ir}} - \mu_i \varphi_r]_{n_r}^+ \\ &\geq \sum_i \frac{1}{\mu_i} \sum_{r \in \mathcal{R}_i} \hat{q}_r (\lambda_i \mathbf{1}_{A_{ir}} - \mu_i \varphi_r), \end{aligned}$$

where in the last step we used the fact that the term inside the projection is negative whenever the projection is active. It follows that:

$$\begin{aligned} \dot{V} &\geq \sum_i \sum_{r \in \mathcal{R}_i} \hat{q}_r (\rho^i \mathbf{1}_{A_{ir}} - \varphi_r) \\ &= \sum_i \sum_{r \in \mathcal{R}_i} \sum_l R_{lr} \hat{p}_l \rho^i \mathbf{1}_{A_{ir}} - \sum_i \sum_{r \in \mathcal{R}_i} \sum_l R_{lr} \hat{p}_l \varphi_r \\ &\geq \sum_i \sum_{r \in \mathcal{R}_i} \sum_l \hat{p}_l R_{lr} \rho^i \mathbf{1}_{A_{ir}} - \sum_l \hat{p}_l c_l. \end{aligned}$$

In the last step, we used the fact that $\sum_r R_{lr} \varphi_r \leq c_l$ due to the resource allocation being feasible. Regrouping the terms we arrive at:

$$\dot{V} \geq \sum_l \hat{p}_l \left(\sum_i \sum_{r \in \mathcal{R}_i} R_{lr} \rho^i \mathbf{1}_{A_{ir}} - c_l \right) \geq a + \epsilon > 0,$$

since $z_l = \sum_r R_{lr} \rho^i \mathbf{1}_{A_{ir}} - c_l \in \mathcal{S}$ by the definition of routing policy. We conclude that V is strictly increasing along the trajectories, and being a linear function of the state, the number of connections would grow without bounds. Thus (20) is necessary for stability. ■

The above Theorem remains valid if we change the policy $\mathbf{1}_{A_{ir}}$ by a random splitting policy, possibly dependent on the state $p_r(n)$, since $p_r(n)$ must also be in the set Δ^i for all n . We therefore have shown that if traffic cannot be split among the routes such that each link is below its capacity on average, then the system cannot be stabilized under any routing policy.

The analogue to the sufficient stability condition in this case is:

$$\forall i, \exists \alpha^i \in \Delta^i : \sum_l R_{lr} \alpha_r^i \rho^i < c_l \quad \forall l, \quad (22)$$

which is the strict version of (20). The following proposition is direct from the single-path stability results of [1], [7]:

Proposition 4: If (22) holds, then there exists a routing policy that stabilizes the system.

Proof: If such a choice of α_r^i exists, then the random splitting policy that sends an incoming connection on route r with probability α_r^i stabilizes the system. This is because the system is equivalent to a single path process with arrival rates $\lambda_i \alpha_r^i$ due to the Poisson thinning property, and condition (22) is the stability condition of the single path case. ■

The above shows that the stability region of this system is characterized completely. However, the random splitting policy mentioned in the proof of Proposition 4 is not useful in a network environment, since it is not decentralized. Each user must know in advance the average loads of the whole system in order to select the weights α_r^i to fulfill (22).

B. A decentralized routing policy

In a multi-path setting in which each user may choose among a set of routes, it is natural to try to balance the loads by using the network congestion prices measured on each route. A simple feedback policy for routing is, at each arrival, choose the route with the cheapest price. In our previous notation this amounts to taking:

$$A_{ir} = \left\{ n : q_r(n) = \min_{r' \in \mathcal{R}_i} q_{r'}(n) \right\}. \quad (23)$$

Implicit in the above equation is some rule for breaking ties when there are multiple routes with minimum price. Since congestion price is inversely related with connection rate, this is equivalent to routing new flows to the path with the best current rate for individual connections. Note also that this is a suitable generalization of the Join the Shortest Queue policy [12]: in fact, in the case of parallel links of equal capacities, the identification is exact. We shall see in Section VII that this closed loop policy does not suffer from the problems of simultaneously using all paths, which can lead to a loss in the stability region, as analyzed in [20].

We shall investigate the stability of this policy under condition (22). We need the following Proposition, proved in the Appendix:

Proposition 5: Given $n = (n_r) \in \mathbb{R}_+^N$, let $\varphi_r(n)$ and $q_r(n)$ be the corresponding rate allocation and route prices from the Congestion Control Problem 1. Choose also α^i that satisfies (22). Then there exists $\delta > 0$ such that:

$$\sum_r q_r(n) \alpha_r^i \rho^i - \sum_{r: n_r > 0} q_r(n) \varphi_r(n) \leq -\delta \sum_r q_r(n). \quad (24)$$

The previous bound is similar to the one used in [1] to prove stability in the single path scenario, but with the gradient evaluated at the optimum, instead of another feasible allocation. We now apply this bound to obtain the following characterization of the routing policy:

Theorem 5: Suppose (22) holds. Then under the dynamics (19) with the routing policy given by (23) we have:

$$\sum_r \frac{q_r \dot{n}_r}{\mu_i} \leq -\delta \sum_r q_r \quad (25)$$

Proof: We have that:

$$\begin{aligned} \sum_r \frac{q_r \dot{n}_r}{\mu_i} &= \sum_r q_r [\rho^i \mathbf{1}_{A_{ir}} - \varphi_r]_{n_r}^+ \\ &\leq \sum_r q_r \rho^i \mathbf{1}_{A_{ir}} - \sum_{r:n_r>0} q_r \varphi_r, \end{aligned}$$

where the inequality is trivial if the projection is not active. If the projection is active, $n_r = 0$ and thus $q_r \varphi_r = 0$ by (28) so the corresponding negative term can be omitted.

Regrouping the above in each user we get:

$$\begin{aligned} \sum_r \frac{q_r \dot{n}_r}{\mu_i} &\leq \sum_i \rho^i \sum_{r \in \mathcal{R}_i} q_r \mathbf{1}_{A_{ir}} - \sum_{r:n_r>0} q_r \varphi_r \\ &= \sum_i \rho^i \min_{r \in \mathcal{R}_i} q_r - \sum_{r:n_r>0} q_r \varphi_r \\ &\leq \sum_i \rho^i \sum_{r \in \mathcal{R}_i} q_r \alpha_r^i - \sum_{r:n_r>0} q_r \varphi_r, \end{aligned}$$

where we have used the definition of the routing policy and the fact that the minimum price can be upper bounded by any convex combination of the prices. Applying Proposition 5 we complete the proof. ■

It remains to see if we can use the inequality (25) to establish asymptotic stability of the fluid dynamics through a Lyapunov argument. Although it is tempting to postulate a Lyapunov function analogous to (21), there is an important difference: the q_r factor of (25) is a function of the state. This implies that the left hand side of (25) may not be integrated easily to get a Lyapunov function for the state space whose derivative along the trajectories yields the desired result.

We focus on a special case where we can give an affirmative answer. Assume that the network is composed by a set of parallel bottleneck links. Each user in this network has a set of routes established in any subset of the links. Moreover, assume that all users have identical α -fair utilities denoted by $U(\cdot)$ and file sizes are equal for each user, so we can take without loss of generality $\mu_i = 1$.

In such a network, the resource allocation of Problem 1 can be explicitly computed as a function of the current number of flows n_r . In particular, all flows through bottleneck l face the same congestion price $q_r = p_l$, and as they have the same utility, they will get the same rate, given by:

$$x_r(n) = \frac{c_l}{\sum_{r' \in l} n_{r'}}.$$

By using that $q_r(n) = U'(x_r(n))$ we have:

$$q_r(n) = x_r(n)^{-\alpha} = \left(\frac{\sum_{r' \in l} n_{r'}}{c_l} \right)^\alpha,$$

with l such that $r \in l$.

Since $q_r = p_l$ if $r \in l$ the link prices should have the form:

$$p_l(n) = \left(\frac{\sum_{r:r \in l} n_r}{c_l} \right)^\alpha.$$

We now state the stability result for this type of network, and defer the proof to the Appendix:

Proposition 6: For the network of parallel links under consideration, let the system be given by (19) with the routing policy (23). Under the stability condition (22), the state n converges to 0 in finite time.

We recall that (c.f. [34, Corollary 9.2]) finite-time convergence to zero of the fluid model implies the stability (ergodicity) of the corresponding Markov chain model.

C. Combining admission control and routing

Let us analyze now the possibility of extending the admission control rules derived in Section V to the connection routing setting. Recall that, in order to perform admission control, we associate with each user a utility function $U_i(\varphi^i)$ with φ^i being the total rate the user gets from the network. Admission control over a route was performed by comparing the marginal utility with the route price. In the new setting, the end user may choose among several routes, and thus the natural way to merge the results of Section V with the connection level routing is the following combined law:

Admit new connection if $\min_{r \in \mathcal{R}_i} q_r < U'_i(\varphi^i)$

If admitted: route connection through the cheapest path

The network dynamics in this case converges to 0 whenever the stability condition (22) is met. In the overload case, it can be shown that the equilibrium is in fact the solution of:

Problem 5 (Saturated User Welfare with Routing):

$$\max_{\varphi_r} \sum_i U_i(\varphi^i)$$

subject to $R\varphi \leq c$ and $\varphi^i \leq \rho^i$ for each i .

The above optimization problem is similar to Problem 3 but the constraint $\varphi^i \leq \rho^i$ is imposed on the aggregate rate of each user and its total average load.

VII. IMPLEMENTATION AND SIMULATIONS

In this section we discuss practical implementation issues and investigate the performance of the policies developed through simulations. We do so in several scenarios: In the first one we consider the case where users cooperate controlling their number of active connections in a proactive way, to achieve a fair resource allocation according to the User Welfare Problem 2. Then we move on to explore the behavior of utility based admission control on an overloaded network, showing that it imposes the desired notion of fairness. Finally we validate the connection level routing policy proposed in Section VI-B in an example where uncoordinated control is known to reduce the stability condition.

A. Scenario 1: Controlling the number of connections.

We implemented a packet level simulation of the control law (10) in the network simulator ns2 [28]. We have two users that download data from 3 servers, with the topologies and link capacities depicted in Fig. 2. In order to introduce an imbalance between users, routes have different round trip

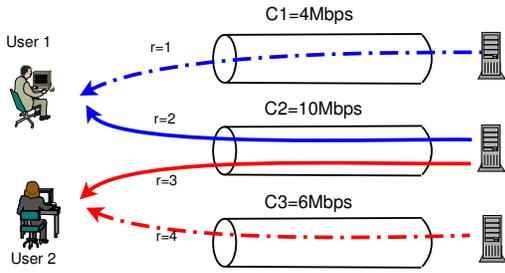


Fig. 2. Topology simulated in Scenario 1.

times. Each user then begins with a single TCP connection per route, governed by TCP/Newreno. With this choice, the congestion price q_r on route r is the packet loss probability along that route, and this is measured by the users counting the number of retransmitted packets within connections.

The users then maintain a variable `targetN(r)` for each route, which is the target number of connections. This variable is updated periodically by measuring the current values of φ^i and q_r and integrating (10). For this particular case, we choose $U_1(\varphi) = U_2(\varphi) = \log(\varphi)$, which will give, in equilibrium, the proportionally fair allocation $\varphi^{1*} = \varphi^{2*} = 10Mbps$.

Each second, a user chooses whether to open or close a connection on route r by comparing the current number of connections along that route with the corresponding `targetN(r)` rounded to the nearest integer. Results in Fig. 3 show that the number of connections on each route tracks an equilibrium value, approximating the dynamics (10). The corresponding total rate evolves reaching the desired allocation. This is achieved by splitting unequally the shared link, but in a decentralized way. A similar scenario, but with uncooperative users, was presented in [11]; in this case admission control can be applied, leading to the same allocation.

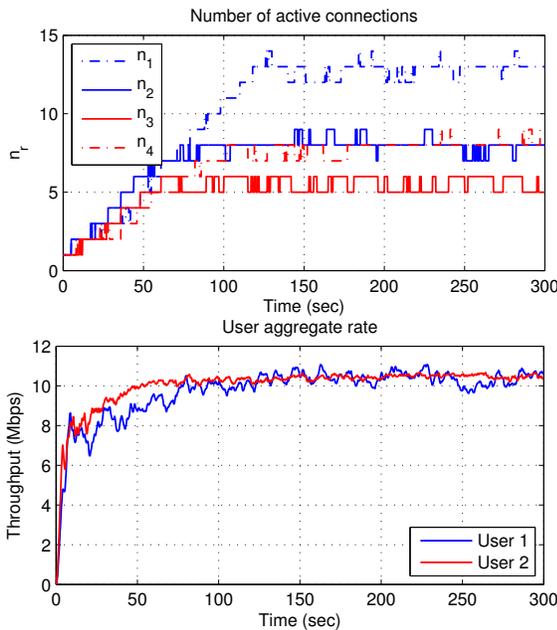


Fig. 3. Results for Scenario 1.

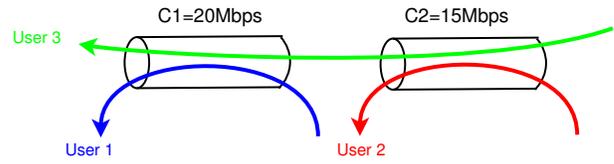


Fig. 4. Topology for Scenario 2.

B. Scenario 2: Fairness via admission control.

In this case we simulated the linear network topology of Fig. 4, with single-path users that generate connections according to a Poisson process as in Section V. The network applies admission control using the rule (12), with utilities chosen from the α fair family with $\alpha = 5$ to approximate max-min fairness. Simulations are performed in ns2 and individual connections are again controlled by TCP/Newreno. The entry router is in charge of measuring loss probability along the routes. This is done in this case by sniffing the connections, although other approaches such as Re-ECN [3] could be used.

The max-min allocation for this network is $\varphi^{1*} = 12.5Mbps$ and $\varphi^{2*} = \varphi^{3*} = 7.5Mbps$. In the first simulation, the user loads are such that the network is overloaded, with each user load ρ^i being greater than its fair share. The results show that the users are admitted $n_1 \approx 8$, $n_2 \approx 10$ and $n_3 \approx 20$ connections in equilibrium, with total rates according to the first graph in Fig. 5. The max-min allocation is approximately achieved. Observe that, despite having the same equilibrium rate than user 2, user 3 is allowed more connections into the network because its RTT is higher and thus its connections are slower.

In the second situation, we changed the load of user 2 to $\rho^2 = 4Mbps$ (below its fair share), and it is therefore saturated

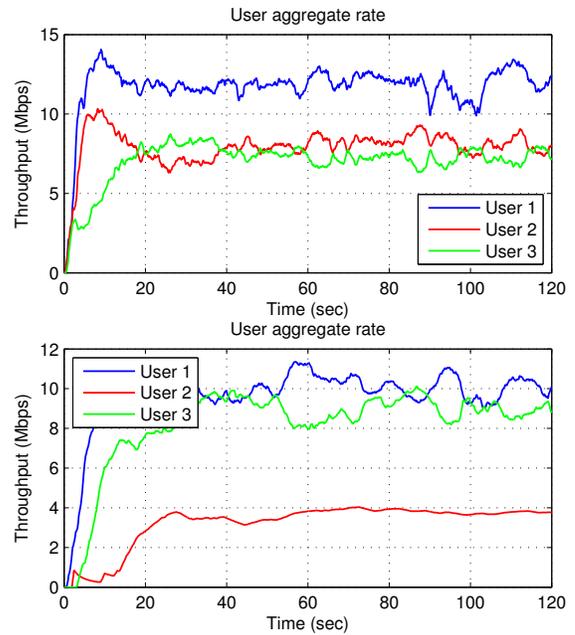


Fig. 5. Results for Scenario 2: fully overloaded case (above) and when one source is below its share (below).

as in Problem 3. The resulting rates are shown in the second graph in Fig. 5. Here admission control is protecting user 2 by allowing its share of $4Mbps$, and reallocating the remaining capacity as in Problem 3 to $\varphi^{1*} = \varphi^{3*} = 10Mbps$, $\varphi^{2*} = 4Mbps$.

C. Scenario 3: Connection level routing.

We now analyze the behavior of the decentralized routing policy for connections presented in Section VI-B. We do so in an example first identified by [20], which corresponds to a network with the following routing matrix:

$$R = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 \end{bmatrix}.$$

The network has 3 links and 3 users, each user with two available paths: a one-hop or a two-hop route. Links are of unit capacity. For the case of symmetric loads ($\rho^i = \rho$ for users $i = 1, 2, 3$), it is easily checked that the stability condition of (22) implies $\rho < 1$, achievable by applying each user's load only in the one-hop path. To satisfy this allocation at the transport layer, a multi-path TCP protocol that coordinates path rates would be required; if instead users open uncoordinated TCP flows in all their routes, it is shown in [20] that the stability region is reduced to $\rho^i < \rho^* = \frac{1+2^{-1/\alpha}}{1+2^{1-1/\alpha}}$ for α -fair TCP, in particular $\rho^* = 1/\sqrt{2} < 1$ for the case $\alpha = 2$.

In our proposal, each connection remains single path and there is no rate coordination between them, but connections are routed to cheapest paths; thus we are able to stabilize the full region $\rho < 1$. In Fig. 6 we show simulation results for a stochastic traffic with loads $\rho^i = 0.8 > \rho^*$, starting from an initial condition of $n_r = 20$ on every route. We can see that the number of connections on long routes is decreasing, and eventually the system converges to a stable behavior. We also show the fluid model trajectory according to dynamics (19) with the rule (23). The fluid model reaches 0 in finite time, consistent with the point at which the stochastic simulation reaches steady state.

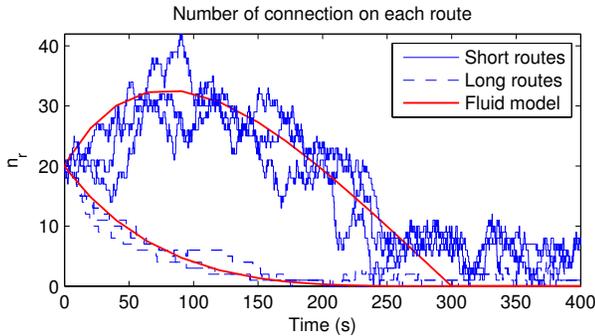


Fig. 6. Results for Scenario 3 and comparison with the fluid model.

VIII. CONCLUSIONS AND FUTURE WORK

In this paper, we analyzed resource allocation in networks from a connection-level perspective, with the intention to

bridge the gap between classical NUM applied to congestion control and the user-centric perspective. New notions of fairness appear, as user utilities are evaluated on aggregates of traffic, which can model different interesting situations. We showed how the control of the number of connections can be used to impose these new notions of fairness, and how the users can cooperate in order to drive the network to a fair equilibrium. Moreover, we showed how admission control and routing based on typical congestion prices can be used to protect the network in overload, and simultaneously impose fairness between its users. Finally, we showed practical implementations of the mechanisms derived in our work, and simulations based on these implementations show that the proposals accomplish their goals.

In future work we plan to address several theoretical and technical questions that are still open. Stability results for admission control, and the stability region of the routing policy proposed are two important theoretical questions. In practical terms, it would be interesting to explore new network implementations based on current congestion notification protocols, that will help make these decentralized admission control mechanisms scalable to large networks.

APPENDIX

Proof of Theorem 1: Consider the map $\Phi : n \mapsto \varphi$ defined by Problem 1. This map is continuous when $n > 0$ [19]. We will also assume that in a neighborhood of the solution, all links are saturated (if there are locally non-saturated links, they can be easily removed from the analysis).

In this case, the KKT conditions of Problem 1 imply:

$$\begin{aligned} U'_{TCP_r} \left(\frac{\varphi_r}{n_r} \right) &= q_r \quad \forall r, \\ \sum_r R_{lr} \varphi_r &= c_l \quad \forall l. \end{aligned}$$

From the first group of equations we have that:

$$\varphi_r = n_r f_{TCP_r}(q_r) \quad \forall r, \quad (26)$$

and substituting in the link constraints we have that the optimal link prices must satisfy:

$$\mathcal{F}(n, p) = R \text{diag}(n) f_{TCP}(R^T p) - c = 0.$$

Here, $\text{diag}(n)$ denotes a diagonal matrix with the entries of n , c is the vector of link rates and $f_{TCP}(R^T p)$ is the vector of flow rates determined by the demand curve in each route.

By using the Implicit Function Theorem we have that:

$$\frac{\partial p}{\partial n} = - \left(\frac{\partial \mathcal{F}}{\partial p} \right)^{-1} \left(\frac{\partial \mathcal{F}}{\partial n} \right).$$

Define now the following matrices:

$$\begin{aligned} N &= \text{diag}(n), \quad F = \text{diag}(f_{TCP}(R^T p)), \\ F' &= \text{diag}(f'_{TCP}(R^T p)). \end{aligned}$$

The diagonal entries of N and F are strictly positive, and the diagonal entries of F' are negative, since we assume the links

are saturated and therefore the prices involved are positive. After some calculations we arrive to:

$$\frac{\partial \mathcal{F}}{\partial p} = RNFR^T, \quad \frac{\partial \mathcal{F}}{\partial n} = RF,$$

and thus:

$$\frac{\partial p}{\partial n} = -(RNFR^T)^{-1}(RF).$$

Note that the first matrix is invertible since R has full row rank and the diagonal matrix has definite sign.

We now turn to calculating $\frac{\partial \varphi}{\partial n}$. From equation (26) we have:

$$\frac{\partial \varphi}{\partial n} = F + NF'R^T \frac{\partial p}{\partial n} = (I - NF'R^T(RNF'R^T)^{-1}R)F,$$

where I is the identity matrix. We would like to prove that the diagonal terms of this matrix are non-negative. Since F is a diagonal matrix with positive entries, we can reduce the problem to proving that the matrix $M = I - DR^T(RDR^T)^{-1}R$ has positive diagonal entries, where $D = -NF'$ is also a diagonal matrix with strictly positive entries, so $\sqrt{D} = \text{diag}(\sqrt{d_{ii}})$ is well defined and:

$$\sqrt{D}^{-1}M\sqrt{D} = I - \sqrt{D}R^T(RDR^T)^{-1}R\sqrt{D},$$

which is of the form $I - A^T(AA^T)^{-1}A$. This last matrix is symmetric and verifies $(I - A^T(AA^T)^{-1}A)^2 = I - A^T(AA^T)^{-1}A$. Thus, it is a self-adjoint projection matrix, and therefore is positive semidefinite. From this we conclude that the diagonal entries of $\sqrt{D}^{-1}M\sqrt{D}$ are non-negative, and since the diagonal entries of M are not altered by this transformation, M has non negative diagonal entries, which concludes the proof. ■

Proof of Theorem 2: Consider the Lyapunov function in (7). Differentiating along the trajectories we get:

$$\begin{aligned} \dot{V} = & \sum_r (\varphi_r - \varphi_r^*) \left[U'_i \left(\sum_{r \in i} \varphi_r \right) - q_r - \nu \dot{q}_r \right]_{\varphi_r}^+ \\ & + \sum_l (p_l - p_l^*) (y_l - c_l)_{p_l}^+ + \sum_l \nu (c_l - y_l^*) \dot{p}_l \end{aligned} \quad (27)$$

Noting that φ_r^* and p_l^* are nonnegative, we can apply the inequality $(v - v_0)(u)_v^+ \leq (v - v_0)u$ to get rid of both positive projections. By inserting the values at equilibrium appropriately we have:

$$\dot{V} \leq \sum_r (\varphi_r - \varphi_r^*) \left(U'_i \left(\sum_{r \in i} \varphi_r \right) - q_r^* \right) \quad (\text{I})$$

$$+ \sum_r (\varphi_r - \varphi_r^*) (q_r^* - q_r) \quad (\text{II})$$

$$- \sum_r \nu (\varphi_r - \varphi_r^*) \dot{q}_r \quad (\text{III})$$

$$+ \sum_l (p_l - p_l^*) (y_l - y_l^*) \quad (\text{IV})$$

$$+ \sum_l (p_l - p_l^*) (y_l^* - c_l) \quad (\text{V})$$

$$+ \sum_l \nu (c_l - y_l^*) \dot{p}_l \quad (\text{VI}).$$

Note that (II) = $-(\varphi - \varphi^*)^T(q - q^*) = -(\varphi - \varphi^*)^T R^T(p - p^*) = -(\text{IV})$ so these terms cancel out. The complementary slackness condition implies (V) ≤ 0 since either $y_l^* = c_l$ and (V) = 0 or $p_l^* = 0$ and (V) ≤ 0 . To bound (I) we associate its terms on each user i to get:

$$\begin{aligned} \sum_{r \in i} (\varphi_r - \varphi_r^*) (U'_i(\varphi^i) - q_r^*) &= \\ &= (\varphi^i - \varphi^{i,*}) (U'_i(\varphi^i) - U'_i(\varphi^{i,*})) \\ &+ \sum_{r \in i} (\varphi_r - \varphi_r^*) (U'_i(\varphi^{i,*}) - q_r^*). \end{aligned}$$

The first term in the right hand side is ≤ 0 since U'_i is increasing. Each of the terms in the sum is also ≤ 0 due to the optimality condition (5), and thus after summing over i we conclude (I) ≤ 0 . The remaining terms can be grouped together and after some manipulations we have:

$$\begin{aligned} (\text{III}) + (\text{VI}) &= -\nu(\varphi - \varphi^*)^T \dot{q} + \nu(c - y^*)^T \dot{p} \\ &= \nu \sum_l \gamma_l (c_l - y_l) (y_l - c_l)_{p_l}^+ \leq 0, \end{aligned}$$

since each term in the last sum is either 0 or $-\gamma_l (y_l - c_l)^2 \leq 0$. We therefore conclude that the function $V(\varphi, p)$ is decreasing along the trajectories. Stability now follows from the LaSalle Invariance Principle [22]. Assume that $\dot{V} \equiv 0$. In particular, the terms (I) and (III) + (VI) must be identically 0 since they are negative semidefinite. Imposing (I) $\equiv 0$ we conclude that $\varphi^i = \varphi^{i,*}$ for all i and $\varphi_r = 0$ for all routes which do not have minimum price. Moreover, $\dot{\varphi}^i = \sum_{r \in i} \dot{\varphi}_r = 0$ since φ^i must be in equilibrium. We also have that:

$$(\text{III}) + (\text{VI}) = \nu \sum_l \gamma_l (c_l - y_l) (y_l - c_l)_{p_l}^+ \equiv 0$$

requires that either $p_l = 0$ or $y_l = c_l$ at all times. Therefore, $\dot{p} = 0$ and p must be in equilibrium. It follows that q_r is in equilibrium, and returning to the dynamics we must have $\dot{\varphi}_r = K_r$ a constant. If $\dot{\varphi}_r = K_r > 0$, it would mean that $\varphi_r \rightarrow \infty$ implying that $y_l \leq c_l$ is violated at some link, contradicting $\dot{p}_l \equiv 0$. Therefore, $K_r \leq 0$ and since $\sum_r \dot{\varphi}_r = 0$, we must have $K_r = 0$. Therefore φ_r is in equilibrium. We conclude that in order to have $\dot{V} \equiv 0$ the system must be in a point that satisfies the KKT conditions (5), and therefore the system will converge to an optimal allocation for Problem 2. Since V is radially unbounded, the convergence holds globally. ■

Remark 3: In the case of a state-dependent gain $k_r \varphi_r > 0$, as required for the dynamics (8), we can change the terms depending on φ_r in (7) to $\int_{\varphi_r^*}^{\varphi_r} \frac{(u - \varphi_r^*)}{k_r u} du$, and equation (27) follows. For further details we refer to [9].

Proof of Theorem 3: The proof is very similar to the fluid limit result from [19], but with additional considerations for the admission control term. We shall use the following stochastic representation of the process $n^L(t)$, in terms of standard Poisson processes with a time scaling. Consider $\{E_r(t)\}$ and $\{S_r(t)\}$ to be a family of independent Poisson processes of intensities λ_r and μ_r respectively. Consider also

the following processes:

$$\begin{aligned}\tau_r^L(t) &= \int_0^t f_\epsilon(U_r^L(\varphi_r^L(n^L(s))) - q_r^L(n^L(s))) ds, \\ T_r^L(t) &= \int_0^t \varphi_r^L(n^L(s)) ds\end{aligned}$$

Here τ_r tracks the amount of time the admission condition is satisfied, and T_r tracks the amount of service provided to route r . The Markov chain evolution (15) of $n^L(t)$ starting at $n^L(0)$ can be written as:

$$n_r^L(t) = n_r^L(0) + E_r(L\tau_r^L(t)) - S_r(T_r^L(t)),$$

where the term $L\tau_r^L(t)$ comes from the fact that the arrival rate of n^L is $L\lambda_r$.

Define $\bar{T}_r^L = T_r^L/L$. Now for each r and $t \geq 0$:

$$\begin{aligned}\bar{T}_r^L(t) &= \frac{1}{L} \int_0^t \varphi_r^L(n^L(s)) ds = \frac{1}{L} \int_0^t \varphi_r^L(L\bar{n}^L(s)) ds \\ &= \frac{1}{L} \int_0^t L\varphi_r(\bar{n}^L(s)) ds = \int_0^t \varphi_r(\bar{n}^L(s)) ds,\end{aligned}$$

where we have used Lemma 1 for φ_r^L . We also have:

$$\begin{aligned}\tau_r^L(t) &= \int_0^t f_\epsilon(U_r^L(\varphi_r^L(n^L(s))) - q_r^L(n^L(s))) ds \\ &= \int_0^t f_\epsilon(U_r^L(L\varphi_r(\bar{n}^L(s))) - q_r(\bar{n}^L(s))) ds \\ &= \int_0^t f_\epsilon(U_r'(\varphi_r(\bar{n}^L(s))) - q_r(\bar{n}^L(s))) ds.\end{aligned}$$

Therefore, the process \bar{n}^L satisfies the following:

$$\begin{aligned}\bar{n}_r^L(t) &= \frac{n_r^L(0)}{L} + \frac{1}{L} E_r(L\tau_r^L(t)) - \frac{1}{L} S_r(L\bar{T}_r^L(t)), \\ \bar{T}_r^L(t) &= \int_0^t \varphi_r(\bar{n}^L(s)) ds, \\ \tau_r^L(t) &= \int_0^t f_\epsilon(U_r'(\varphi_r(\bar{n}^L(s))) - q_r(\bar{n}^L(s))) ds.\end{aligned}$$

The conclusion now follows from the same arguments in [19] applying the hypothesis for $n^L(0)$ and the functional law of large numbers for the processes E_r and S_r , namely, $\frac{1}{L}E_r(Lu) \rightarrow \lambda_r u$ and $\frac{1}{L}S_r(Lu) \rightarrow \mu_r u$. Note that the functions \bar{T}_r are Lipschitz since they are the integral of a bounded function (φ_r is bounded by the maximum capacity in the network). Also, τ_r^L is Lipschitz because f_ϵ is bounded by 1. We conclude that any limit point of \bar{n}^L must satisfy:

$$\begin{aligned}\bar{n}_r &= \bar{n}_r(0) + \lambda_r \int_0^t f_\epsilon(U_r'(\varphi_r(\bar{n}(s))) - q_r(\bar{n}(s))) ds \\ &\quad - \mu_r \int_0^t \varphi_r(\bar{n}(s)) ds.\end{aligned}$$

Differentiating the above equation at any regular point of the limit, we get the desired result. ■

Proof of Proposition 5: The proof of the Proposition is based on the following result, proved in [9], which deals with the quantities $q_r(n)\varphi_r(n)$ when $n_r \rightarrow 0$. Let \bar{n} be such that $\bar{n}_r = 0$ for some r . Then we have:

$$\lim_{n \rightarrow \bar{n}, n_r > 0} q_r(n)\varphi_r(n) = 0 \quad (28)$$

Since the inequality in (22) is strict, we can choose $\delta > 0$ such that $\psi_r := \alpha_r^i \rho^i + \delta$, $\forall r$ is such that $R\psi \leq c$.

Consider now $n > 0$. Since $\varphi(n)$ is the optimal allocation, and ψ is another feasible allocation for the convex Problem 1, the inner product must satisfy:

$$\nabla_\varphi \left(\sum_r n_r U_{TCP_r}(\varphi_r/n_r) \right) (\psi - \varphi) \leq 0.$$

Otherwise, one could improve the utility of solution φ by moving it in the direction of ψ .

Applying the above condition to $\psi_r = \alpha_r^i \rho^i + \delta$ and using

$$\frac{\partial (\sum_r n_r U_{TCP_r}(\varphi_r/n_r))}{\partial \varphi_r} = U'_{TCP_r} \left(\frac{\varphi_r}{n_r} \right) = q_r(n),$$

we conclude that:

$$\sum_r q_r(n) (\alpha_r^i \rho^i + \delta - \varphi_r(n)) \leq 0.$$

The above can be rewritten as

$$\sum_r q_r(n) \alpha_r^i \rho^i - \sum_r q_r(n) \varphi_r(n) \leq -\delta \sum_r q_r(n),$$

which proves the result for $n > 0$ component-wise.

If now n is such that $n_r = 0$ for some r , we can take limits from points inside the orthant $n > 0$ and use (28) to obtain the result. ■

Proof: Consider the candidate Lyapunov function:

$$V(n) = \frac{1}{\alpha + 1} \sum_l \frac{1}{c_l^\alpha} \left(\sum_{r:r \in l} n_r \right)^{\alpha+1}. \quad (29)$$

The above function is continuous and non-negative in the state space, radially unbounded and is only 0 at the equilibrium $n = 0$. Its derivative along the trajectories verifies:

$$\dot{V} = \sum_l \frac{(\sum_{r:r \in l} n_r)^\alpha}{c_l^\alpha} \left(\sum_{r:r \in l} \dot{n}_r \right) = \sum_r q_r \dot{n}_r. \quad (30)$$

Invoking Theorem 5 we conclude that $\dot{V} \leq 0$ in the state space, and it is only 0 when all prices are 0 which can only happen at the origin. This implies asymptotic stability of the fluid dynamics.

To obtain finite time convergence note that V verifies:

$$V(n) = \frac{1}{\alpha + 1} \sum_l p_l^{\frac{\alpha+1}{\alpha}} c_l \leq (K \max_l p_l)^{\frac{\alpha+1}{\alpha}},$$

where $K > 0$ is an appropriate constant.

We can obtain the following bound:

$$V^{\frac{\alpha}{\alpha+1}} \leq K \max_l p_l \leq K \sum_l p_l \leq K \sum_r q_r,$$

and thus $-\delta \sum_r q_r \leq -\frac{\delta}{K} V^{\frac{\alpha}{\alpha+1}}$. Combining this with the result from Theorem 5 we get $\dot{V} \leq -\frac{\delta}{K} V^{\frac{\alpha}{\alpha+1}}$.

Integrating the above inequality yields:

$$V(t)^{\frac{1}{1+\alpha}} \leq V(0)^{\frac{1}{1+\alpha}} - \frac{\delta}{K(1+\alpha)} t,$$

and we conclude that V , and therefore n , reach 0 in finite time, proportional to $V(0)^{\frac{1}{1+\alpha}}$. ■

REFERENCES

- [1] T. Bonald and L. Massoulié, "Impact of fairness on internet performance," in *ACM SIGMETRICS/Performance*, 2001, pp. 82–91.
- [2] S. Boyd and L. Vanderberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [3] B. Briscoe, "Flow rate fairness: Dismantling a religion," *Computer Communication Review*, vol. 37, no. 2, pp. 63–74, 2006.
- [4] M. Chen and A. Zakhor, "Flow control over wireless network and application layer implementation," in *IEEE Infocom*, Apr. 2006.
- [5] —, "Rate control for streaming video over wireless," in *IEEE Infocom*, Mar. 2004.
- [6] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," in *Proceedings of the IEEE*, vol. 95, no. 1, Jan 2007, pp. 255–312.
- [7] 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*, 1999, pp. 802–810.
- [8] D. Feijer and F. Paganini, "Stability of primal-dual gradient dynamics and applications to network optimization," *Automatica*, vol. 46, no. 12, pp. 1974–1981, 2010.
- [9] A. Ferragut, "Resource allocation in networks from a connection level perspective," Ph.D. dissertation, Universidad de la República, Uruguay, 2011.
- [10] A. Ferragut and F. Paganini, "Achieving network stability and user fairness through admission control of TCP connections," in *Conference on Information Sciences and Systems*, 2008.
- [11] —, "User-centric network fairness through connection-level control," in *IEEE Infocom*, 2010.
- [12] R. D. Foley and D. R. McDonald, "Join the shortest queue: stability and exact asymptotics," *Annals of Applied Probability*, vol. 11, no. 3, pp. 569–607, 2001.
- [13] T. J. Hacker, B. D. Athey, and B. Noble, "The end-to-end performance effects of parallel TCP sockets on a lossy wide-area network," in *Proceedings of the 16th International Parallel and Distributed Processing Symposium (IPDPS 2002)*, Apr. 2002.
- [14] H. Han, S. Shakkotai, C. Hollot, R. Srikant, and D. Towsley, "Multipath TCP: A joint congestion control and routing scheme to exploit path diversity in the internet," *IEEE/ACM Trans. on Networking*, vol. 14, no. 6, pp. 1260–1271, 2006.
- [15] M. Jonckheere and J. Virtamo, "Optimal insensitive routing and bandwidth sharing in simple data networks," in *ACM SIGMETRICS*, Jun. 2005.
- [16] F. Kelly, P. Key, and S. Zachary, "Distributed admission control," *IEEE Journal on Selected Areas in Communication*, vol. 18, no. 12, pp. 2617–2628, 2000.
- [17] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 39, pp. 237–252, 1998.
- [18] F. Kelly and T. Voice, "Stability of end-to-end algorithms for joint routing and rate control," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 2, pp. 5–12, 2005.
- [19] F. P. Kelly and R. Williams, "Fluid model for a network operating under a fair bandwidth sharing policy," *Annals of Applied Probability*, vol. 14, no. 3, pp. 1055–1083, 2004.
- [20] P. Key and L. Massoulié, "Fluid models of integrated traffic and multipath routing," *Queueing Systems*, vol. 53, no. 1–2, pp. 85–98, 2006.
- [21] P. Key, L. Massoulié, and D. Towsley, "Path selection and multipath congestion control," in *IEEE Infocom*, Mar. 2007.
- [22] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice-Hall, 1996.
- [23] R. Khalili, N. Gast, M. Popovic, U. Upadhyay, and J.-Y. L. Boudec, "MPTCP is not Pareto-optimal: performance issues and a possible solution," in *Proceedings of the ACM CoNEXT'12*, Dec. 2012.
- [24] X. Lin, N. Shroff, and R. Srikant, "On the connection-level stability of congestion-controlled communication networks," *IEEE Trans. on Information Theory*, vol. 54, no. 5, pp. 2317–2338, 2007.
- [25] X. Lin and N. B. Shroff, "Utility maximization for communication networks with multi-path routing," *IEEE/ACM Trans. on Automatic Control*, vol. 51, no. 5, pp. 766–781, 2006.
- [26] L. Massoulié, "Structural properties of proportional fairness: stability and insensitivity," *Annals of Applied Probability*, vol. 17, no. 3, pp. 809–839, 2007.
- [27] L. Massoulié and J. Roberts, "Arguments in favour of admission control of TCP flows," in *International Teletraffic Conference 16*, Jun. 1999.
- [28] S. McCanne and S. Floyd, "ns Network Simulator," <http://www.isi.edu/nsnam/ns/>, 2000.
- [29] J. Mo and J. Walrand, "Fair end-to-end window based congestion control," *IEEE/ACM Trans. on Networking*, vol. 8, no. 5, pp. 556–567, Oct. 2000.
- [30] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling tcp throughput: a simple model and its empirical validation," *ACM SIGCOMM Computer Communication Review*, vol. 28, no. 4, pp. 303–314, 1998.
- [31] F. Paganini and E. Mallada, "A unified approach to congestion control and node-based multipath routing," *IEEE/ACM Trans. on Networking*, vol. 17, no. 5, pp. 1413–1426, 2009.
- [32] F. Paganini, K. Tang, A. Ferragut, and L. Andrew, "Stability of networks under general file size distribution and alpha fair bandwidth allocation," *IEEE Trans. on Automatic Control*, vol. 57, no. 3, pp. 579–591, 2012.
- [33] C. Raiciu, M. Handley, and D. Wischik, "Rfc 6356: Coupled congestion control for multi-path transport protocols," Oct. 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6356.txt>
- [34] P. Robert, *Stochastic Networks and Queues*, ser. Stochastic Modelling and Applied Probability Series. New York: Springer-Verlag, 2003.
- [35] R. Srikant, *The Mathematics of Internet Congestion Control*. Boston, MA: Birkhäuser, 2004.
- [36] T. Voice, "Stability of multi-path dual congestion control algorithms," in *IEEE Valuetools'06*, Oct. 2006.
- [37] —, "Stability of multi-path dual congestion control algorithms," *IEEE/ACM Trans. on Networking*, vol. 15, no. 6, pp. 1231–1239, 2007.
- [38] J. Wang, J. Wen, J. Zhang, and Y. Han, "TCP-FIT: An improved TCP congestion control algorithm and its performance," in *IEEE Infocom*, Mar. 2011.
- [39] D. Wischik, C. Raiciu, A. Greenhalgh, and M. Handley, "Design, implementation and evaluation of congestion control for Multipath TCP," in *Proceedings of USENIX NSDI 2011*, Mar. 2011.
- [40] H. Zhang, D. F. Towsley, and W. Gong, "TCP connection game: A study on the selfish behavior of TCP users," in *ICNP'05*, 2005, pp. 301–310.



Andrés Ferragut Andrés Ferragut obtained his PhD degree in Electrical Engineering (Telecommunications) from Universidad de la República, Uruguay (2011). He has held teaching and research positions at the Mathematics and Electrical Engineering Depts. at Univ. de la República (2000 to 2009) and since 2007 at the School of Engineering, Universidad ORT Uruguay, where he is currently affiliated with the Mathematics Applied to Telecommunications (MATE) research group. He also spent two years (2004-2006) of graduate work at Télécom Paris, France. In 2012, Dr. Ferragut was awarded the annual Uruguayan Engineering Academy prize for the best doctoral thesis in Electrical Engineering. His research interests are in mathematical modeling applied to telecommunication networks.



Fernando Paganini Fernando Paganini (M'90–SM'05) received his Ingeniero Electricista and Licenciado en Matemática degrees from Universidad de la República, Montevideo, Uruguay, in 1990, and his M.S. and PhD degrees in Electrical Engineering from the California Institute of Technology, Pasadena, in 1992 and 1996 respectively. His PhD thesis received the 1996 Wilts Prize and the 1996 Clauser Prize at Caltech. From 1996 to 1997 he was a postdoctoral associate at MIT. Between 1997 and 2005 he was on the faculty the Electrical Engineering Department at UCLA, reaching the rank of Associate Professor. Since 2005 he is Professor of Electrical and Telecommunications Engineering at Universidad ORT Uruguay. Dr. Paganini has received the 1995 O. Hugo Schuck Best Paper Award, the 1999 NSF CAREER Award, the 1999 Packard Fellowship, the 2004 George S. Axelby Best Paper Award, and the 2010 Elsevier Scopus Prize. He is a member of the Uruguayan National Academy of Sciences. His research interests are control and networks.