

Utility Max-Min Fair Congestion Control with Time-Varying Delays

Konstantin Miller
Berlin Institute of Technology
Email: miller@tkn.tu-berlin.de

Tobias Harks
Berlin Institute of Technology
Email: harks@math.tu-berlin.de

Abstract—We present a framework for designing delay-independent end-to-end congestion control algorithms, where each end-user may have a different utility function. We only require that utility functions are strictly increasing. In this framework, we design an algorithm that maximizes the minimum utility value in the network, that is, the resulting resource allocation is utility max-min fair. To achieve this, we first extend the congestion control algorithm EMKC proposed by Zhang et al. [1], which aims at max-min fair bandwidth allocation. Our extension (xMKC) allows for *arbitrary* rate allocations in the steady state. We investigate xMKC analytically and prove local asymptotic stability with heterogeneous time-varying feedback delays in multi-link networks and global asymptotic stability with homogeneous time-varying feedback delays in single-link networks. Then, we propose UMKC (Utility Max-Min Fair Kelly Control), which achieves utility max-min fairness in its steady state. Based on the analysis of xMKC, we establish stability results for UMKC in the presence of time-varying feedback delays. Finally, we evaluate the performance of UMKC using NS-2 simulations [2].

I. INTRODUCTION

In the last years, congestion control algorithms for communication networks (including various versions of TCP [3]–[7]) have been interpreted as distributed algorithms at sources and links, which implicitly solve a global optimization problem, see [8]–[11] and references therein. In these models, users are associated with increasing and strictly concave utility functions representing elastic traffic. Canonical distributed congestion control algorithms that maximize aggregate utility are derived by decomposing the optimization problem into subproblems that can be solved by links and sources using only local information. The links communicate a price based on usage measurements and sources adapt the sending rates based on the aggregate price (or maximum price as in [1], [12]). In addition to the strong design goal of having distributed algorithms, asymptotic stability in presence of feedback delays is important for congestion control algorithms in high-speed networks. To further complicate the situation, such delays are by nature *time-varying*, since they are a result of perturbations caused by randomly arriving and departing flows as well as varying queuing delays.

Even though considerable progress has been made in designing stable and distributed congestion control algorithms, two important issues have only partly been addressed.

First, most of the stability results that have been obtained so far deal either with fixed homogeneous delays ($d_i = d$) [8],

[13]–[19], or with fixed heterogeneous delays [20]–[22]. In these works, it is required that flows adapt their sending rates based on feedback delays, which may lead to severe RTT-unfairness as shown in [23]. Based on ideas of [9], [14], [21], Zhang and Loguinov [24] remedied the delay dependency by proposing EMKC, a delay-independent max-min fair congestion control algorithm. They showed that systems with fixed diagonal feedback delays are stable if the Jacobian in the steady state is diagonally similar to a radial matrix (matrix A is radial iff $\|A\|_2 = \rho(A)$) and $\rho(J) < 1$. Still, their results do not cover time-varying delays.

The second issue concerns the design of efficient, fair and stable congestion control algorithms incorporating *inelastic* traffic. Many works on congestion control algorithms using the utility framework are focused on elastic traffic. As shown in [25]–[28], some applications, especially real-time applications, have non-concave utility functions. Only a few works deal with this non-concavity, see for instance [29], [30]. In [31] it was shown that the canonical distributed algorithms may fail to converge to a feasible rate allocation and may lead to instability and congestion. To overcome these problems the authors propose a ‘self regulating’ heuristic for the special type of sigmoidal utilities in combination with a subgradient method to generate prices. In [32], non-concave utility functions are assumed and the conditions for convergence to global optimality of the canonical algorithms are analyzed.

Another line of research dealing with non-convexity is concerned with a different fairness definition rather than maximizing aggregate utility: an equilibrium point should result in roughly equal utility values for different applications. In [33], only mild assumptions on the feasible utility functions are required (non-decreasing, not necessarily continuous). In this approach, the links have to maintain per-flow states in order to allocate bandwidth utility fair. Furthermore, no stability results are given in the presence of communication delays. The work in [29] presented a link algorithm that achieves a utility max-min fair bandwidth allocation, where for each link the utility functions of all flows sharing that link is maintained. In [34], [35] distributed algorithms without per-flow states are presented that converge to a utility max-min fair operating point. They, however, prove stability only under fixed communication delays and single link networks.

Our goal in this work is to contribute to the above two issues. First, we propose xMKC, a generalization of EMKC,

which allows for arbitrary bandwidth allocations in the steady state but requires centralized calculation of control parameters. xMKC is intended to serve as a base for algorithms that achieve desired rate allocations in a fully distributed manner. For xMKC, we prove local asymptotic stability with heterogeneous time-varying feedback delays in single and multi-link networks and global asymptotic stability with homogeneous time-varying feedback delays in single-link networks. Based on xMKC, we develop an algorithm called uMKC (Utility Max-Min Fair Kelly Control), which provides a unifying congestion control algorithm for elastic and real-time traffic and is fully distributed. We prove that uMKC's steady state is utility max-min fair in arbitrary network topologies. Furthermore, based on the results for xMKC, we establish stability results for uMKC in the presence of time-varying feedback delays. Finally, we simulate uMKC using the network simulator NS-2 [2].

II. SYSTEM MODEL

We model a packet switched network by a set \mathcal{L} of unidirectional links with capacities $(C_l, l \in \mathcal{L})$. Here, capacities are not necessarily physical constraints but can be some configured target values of the aggregate load at the links. The set of links is shared by a set \mathcal{S} of sources (we use the terms source, application and flow interchangeably). Each source $i \in \mathcal{S}$ represents an end-to-end connection that involves the subset $\mathcal{L}(i)$ of links. Equivalently, each link l is used by the set $\mathcal{S}(l)$ of sources. In this work, we do not consider multi-path routing. We denote the integral time variable with $t \in \mathbb{Z}_0^+$ and the rate of flow i at time instant t with $x_i(t) \in \mathbb{R}_0^+$. Each application is associated with a utility function $U_i(x_i) : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ that fulfills the conditions

- 1) $U_i(0) = 0, \forall i \in \mathcal{S}$,
- 2) $U_i(x)$ is continuous and strictly increasing $\forall i \in \mathcal{S}$ and $\forall x \in \mathbb{R}_0^+$.

Note that our results are not restricted to strict concave utility functions, which makes them applicable not only for elastic traffic like TCP but also for multimedia applications with non-concave utility functions [25].

For each source $i \in \mathcal{S}$ and each link $l \in \mathcal{L}(i)$, we denote by $d_{i,l}^{\rightarrow} \in \mathbb{Z}_0^+$ the delay on the path from the source i to the link l . Conversely, the value $d_{i,l}^{\leftarrow} \in \mathbb{Z}^+$ denotes the delay of the reverse path that includes the path from the link l to the destination and the path from the destination back to the source. For the round-trip delay we obtain $d_i = d_{i,l}^{\leftarrow} + d_{i,l}^{\rightarrow} = d_{i,l'}^{\leftarrow} + d_{i,l'}^{\rightarrow} \geq 1, \forall l, l' \in \mathcal{L}(i)$. Note that with this notation, a value of 1 represents an undelayed link.

For each link l , the aggregate load at time t is given by $X_l(t) = \sum_{i \in \mathcal{S}(l)} x_i(t - d_{i,l}^{\rightarrow})$. We call a rate allocation $x = (x_i, i \in \mathcal{S})$ *feasible*, if and only if $X_l \leq C_l, \forall l \in \mathcal{L}$.

III. xMKC: A GENERALIZATION OF EMKC

In this section, we generalize the source equation of EMKC in order to achieve arbitrary rate allocations in the steady state.

We also modify the router equation in order to avoid capacity overshoots in the steady state as observed with EMKC. This modification allows us to establish stability results in the presence of time-varying feedback delays, which we present in Section III-B. Note, that in [24], Zhang and Loguinov also established local stability for the generalized source equation. However, their results are only valid for time-invariant feedback delays.

In order to achieve a rate allocation according to a certain allocation strategy or fairness criterion, xMKC requires a centralized calculation of control parameters. Our analysis of xMKC, however, proves useful for deriving stability results for distributed congestion control algorithms. Specifically, in Section IV we will present uMKC, which achieves utility max-min fairness in a distributed manner. We will also show that uMKC can be used to allocate bandwidth in a max-min fair way, since max-min fairness complies with utility max-min fairness, when all users have the utility function $U(x) = x$.

A. Generalization of EMKC's steady state

Our first goal is to modify EMKC to achieve arbitrary steady state rate allocations. To achieve this goal, we modify the source equation of EMKC. We obtain:

$$x_i(t) = x_i(t - d_i) + \kappa [w_i - \eta_i(t) \cdot x_i(t - d_i)]. \quad (1)$$

Here, $\eta_i(t)$ is the maximum congestion measure of all links on the path of flow i : $\eta_i(t) := \max_{l \in \mathcal{L}(i)} p_l(t - d_{i,l}^{\leftarrow})$. The difference between (1) and the sender equation of EMKC is that we allow each sender to have individual values w_i , whereas with EMKC $w_i = w_j, \forall i, j \in \mathcal{S}$. We will show that with a proper choice of w_i and with a modified version of the router equation that we present below we can achieve arbitrary rate allocations in the steady state.

To deal with the capacity overshoot as observed with EMKC, we modify the router equation for the congestion measure $p_l(t)$. We introduce the value $\tilde{C}_l = \nu C_l$ with $\nu < 1$, which we call the *virtual capacity*, and use it instead of the target capacity C_l to calculate $p_l(t)$. In Section III-B, we show that the choice of ν is critical to system's stability and how it has to be chosen to guarantee stability in presence of time-varying feedback delays. The modified router equation is:

$$p_l(t) = \frac{X_l(t) - \tilde{C}_l}{X_l(t)}. \quad (2)$$

The following Proposition describes the equilibrium structure of xMKC for a given parameter vector $w = (w_i, i \in \mathcal{S})$. In the following, we will omit the router index l if considering single-link scenarios.

Proposition 1. *In a single-link network, the equilibrium structure of xMKC, as described by (1), (2), with fixed $\nu \in (0, 1)$, $\kappa \in (0, 1)$, and $w = (w_i, i \in \mathcal{S})$ with $w_i > 0, \forall i \in \mathcal{S}$, has the*

shape

$$\hat{x}_i = \frac{\mu}{\mu - \nu} \cdot w_i, \quad (3)$$

$$\hat{X} = \mu C, \quad (4)$$

where $\mu = \nu + \frac{1}{C} \sum w_i$.

We omit the proof for brevity.¹

From this Proposition we see, how the equilibrium rates \hat{x}_i depend on the parameter vector w . In the following Proposition, we answer the reverse question of how to choose the parameter vector w in order to achieve a given bandwidth allocation z .

Proposition 2. *In a single-link network with fixed $\nu \in (0, 1)$, $\kappa \in (0, 1)$, and given a rate allocation $z = (z_i, i \in \mathcal{S})$ with $Z := \sum z_i = \mu C$ for a $\mu > \nu$, there exists a unique parameter vector $w = (w_i, i \in \mathcal{S})$ such that xMKC, as described by (1), (2), achieves rate allocation z in its steady state. The corresponding parameter vector is given by*

$$w_i = \frac{\mu - \nu}{\mu} z_i, \quad \forall i \in \mathcal{S}. \quad (5)$$

Proof: First, we show that choosing w according to (5) results in the desired rate allocation z . For the sum of the control parameters w_i , we obtain

$$\begin{aligned} \sum w_i &= \frac{\mu - \nu}{\mu} \sum z_i = \frac{\mu - \nu}{\mu} \mu C = (\mu - \nu) C \\ \Leftrightarrow \mu &= \nu + \frac{1}{C} \sum w_j. \end{aligned}$$

Applying Proposition 1, we obtain $\hat{x}_i = z_i$. For the proof of uniqueness, consider another parameter vector $\tilde{w} = (\tilde{w}_i, i \in \mathcal{S})$. Assuming $\sum w_j \neq \sum \tilde{w}_j$, we obtain $\mu' = \nu + \frac{1}{C} \sum \tilde{w}_j \neq \mu$ and therefore with Proposition 1: $\hat{X}' = \mu' C \neq Z$. Therefore we know $\sum \tilde{w}_j = \sum w_j$. Applying once again Proposition 1, we obtain $\hat{x}_i = z_i = \frac{\mu}{\mu - \nu} w_i = \frac{\mu}{\mu - \nu} \tilde{w}_i, \forall i \in \mathcal{S}$, from which follows $w_i = \tilde{w}_i, \forall i \in \mathcal{S}$. ■

Usually, calculation of the desired rate allocation z that satisfies a certain fairness criterion, requires knowledge of the networks parameters. For now, we assume that z is given. In Section IV, we will use these results to reach a utility max-min fair rate allocation in a fully distributed manner.

B. Stability of xMKC

One of the most important characteristics of a congestion control algorithm is its asymptotic stability. An asymptotically stable algorithm is more likely to avoid oscillations in its steady state and to properly respond to perturbations caused by arrival and departure of flows and other transient effects.

1) *Local Stability:* Zhang and Loguinov [24] established a sufficient condition for local asymptotic stability with heterogeneous time-invariant feedback delays for EMKC for a broader class of feedback functions $p(t)$. In real networks, however, feedback delays vary with time. To establish a sufficient condition for stability in presence of time-varying

feedback delays, we use a result obtained by Kaszkurewicz and Bhaya in [36]. The authors show that a non-linear discrete system is locally asymptotically stable in its stationary point in the presence of heterogeneous uniformly-bounded time-varying delays if all eigenvalues of $|J|$ are smaller than 1. Here, J denotes the Jacobian of the system in the stationary point and $|A| := [|a_{ij}|]_{i,j \in \mathcal{S}}$.

Theorem 1. *In a single-link network with heterogeneous time-varying feedback delays xMKC, as described by (1), (2), with an arbitrary parameter vector $w = (w_i, i \in \mathcal{S})$ and fixed $\nu \in (0, 1)$ and $\kappa \in (0, 1)$, is locally asymptotically stable in its steady state if*

$$\sum w_i \geq \nu C. \quad (6)$$

Proof: Using the steady state values from Proposition 1, we can calculate the Jacobian of the undelayed system in the equilibrium point:

$$\begin{aligned} \left. \frac{\partial p(t-1)}{\partial x_i(t-1)} \right|_{ep} &= \frac{\tilde{C}}{\hat{X}^2} = \frac{\nu}{\mu^2 \cdot C}, \\ \left. \frac{\partial x_i(t)}{\partial x_i(t-1)} \right|_{ep} &= 1 - \frac{\kappa}{\mu(\mu - \nu)C} [\nu w_i + (\mu - \nu)^2 C], \\ \left. \frac{\partial x_i(t)}{\partial x_j(t-1)} \right|_{ep} &= -\kappa \frac{\nu w_i}{\mu(\mu - \nu)C}. \end{aligned}$$

Since the last equation does not depend on j , the Jacobian in the steady state has the following shape (with $n := |\mathcal{S}|$):

$$J = \begin{pmatrix} a_1 & b_1 & \dots & b_1 \\ b_2 & a_2 & \dots & b_2 \\ \dots & \dots & \dots & \dots \\ b_n & b_n & \dots & a_n \end{pmatrix},$$

where $a_i = \left. \frac{\partial x_i(t)}{\partial x_i(t-1)} \right|_{ep}$ and $b_i = \left. \frac{\partial x_i(t)}{\partial x_j(t-1)} \right|_{ep}$.

To complete the proof, we bound the eigenvalues of $|J|$ from above by the L^1 matrix norm of $|J|$ and then show that condition (6) implies: $\| |J| \|_1 < 1$. Simple transformations result in: $1 - \kappa \leq a_i < 1$ and $-\kappa < b_i < 0$. Since we demanded that $\kappa \in (0, 1)$, we obtain:

$$|J| = \begin{pmatrix} a_1 & -b_1 & \dots & -b_1 \\ -b_2 & a_2 & \dots & -b_2 \\ \dots & \dots & \dots & \dots \\ -b_n & -b_n & \dots & a_n \end{pmatrix}.$$

The eigenvalues of a matrix are bounded by any of the induced matrix norms. Thus, we obtain for the eigenvalues of $|J|$

$$|\lambda_k| \leq \| |J| \|_1 = \max_{i \in \mathcal{S}} \left[a_i - \sum_{j \neq i} b_j \right], \quad \forall k \in \mathcal{S}.$$

It is now sufficient to show that $a_i - \sum_{j \neq i} b_j < 1, \forall i \in \mathcal{S}$.

¹All proves will appear in the full version of the paper.

To do this, we explicitly calculate the above sum and obtain

$$a_i - \sum_{j \neq i} b_j = 1 + \frac{\kappa}{\mu(\mu - \nu)C} \left[\nu \sum_{j \neq i} w_j - \nu w_i - (\mu - \nu)^2 C \right]$$

$$< 1 + \frac{\kappa}{\mu(\mu - \nu)C} \left[\nu \sum_{j \in \mathcal{S}} w_j - (\mu - \nu)^2 C \right] = 1 + \kappa \frac{2\nu - \mu}{\mu},$$

which is smaller than or equal to 1 if $\nu \leq \frac{\mu}{2}$, which in turn is an implication of condition (6). ■

Remark. Condition (6) can be translated into an equivalent condition on the target load and parameter ν . For example, to ensure local stability of any steady state with an aggregate load $\hat{X} = C$, it is sufficient to set $\nu \leq 0.5$.

2) *Global Stability:* Global stability analysis of non-linear systems is a very challenging issue. Zhang and Loguinov [37] proved global asymptotic stability of EMKC in single-link networks with homogeneous time-invariant feedback delays for the case $w_i = w_j, \forall i, j \in \mathcal{S}$. We present a fairly simpler proof of global stability of xMKC for the general case $w_i \neq w_j$ in the presence of homogeneous time-varying feedback delays, which is an assumption that is more close to the conditions in real networks. For the proof, we will need the following Lemma:

Lemma 1. A sequence given by the recursive equation $x_n = a_{n-1} \cdot x_{n-1} + b_{n-1}$ can be written in the explicit form

$$x_n = x_0 \cdot \prod_{i=0}^{n-1} a_i + \sum_{i=0}^{n-1} b_i \prod_{j=i+1}^{n-1} a_j.$$

Additionally, if $a_n = a$ and $b_n = b, \forall n \in \mathbb{N}$, this expression can be written as

$$x_n = a^n \cdot x_0 + b \cdot \frac{1 - a^n}{1 - a}.$$

With this Lemma we can now prove the following Theorem:

Theorem 2. In a single-link network with homogeneous time-varying feedback delays xMKC, as described by (1), (2), with fixed $\nu \in (0, 1)$ and $\kappa \in (0, 1)$ and with an arbitrary parameter vector $w = (w_i, i \in \mathcal{S})$, is globally asymptotically stable.

Proof: We establish the proof in two steps. First, we will show that the aggregate load converges to its steady state value for arbitrary initial conditions. This step is analogous to [1], where it is done for $w_i = w_j, \forall i, j \in \mathcal{S}$, and for fixed delays. Then, we will show that convergence of the total load implies convergence of the individual sending rates.

Since all users are assumed to have the same delays, we can write $d_{i,l}^+(t) = d_{j,l}^+(t) = d^+(t), \forall i, j \in \mathcal{S}(l)$, and $d_{i,l}^-(t) = d_{j,l}^-(t) = d^-(t), \forall i, j \in \mathcal{S}(l)$. The sender equation then becomes

$$x_i(t) = x_i(t - d(t)) + \kappa [w_i - p(t - d^-(t)) \cdot x_i(t - d(t))].$$

We denote by $t_k \in \mathbb{Z}_0^+$ the sequence of time instants that are on the time-scale of the feedback delays. That is:

$t_k = t_{k-1} + d(t_{k-1}), t_0 \in [0, d(0))$. Writing $x_i^{(k)}$ for $x_i(t_k)$ and $p^{(k-1)}$ for $p(t_k - d^-(t_k))$, we can rewrite the sender equation as

$$x_i^{(k)} = x_i^{(k-1)} + \kappa [w_i - p^{(k-1)} \cdot x_i^{(k-1)}].$$

Substituting the congestion measure (2) in the above equation and taking the sum over i , we obtain with $X^{(k)} := X(t_k + d^+(t_k))$:

$$\sum x_i^{(k)} = X^{(k)} = (1 - \kappa)X^{(k-1)} + \kappa\mu C,$$

where $\mu = \nu + \frac{1}{C} \sum w_i$.

Applying Lemma 1, we can write this expression in an explicit form:

$$X^{(k)} = (1 - \kappa)^k X^{(0)} + (1 - (1 - \kappa)^k) \mu C.$$

This sequence converges to μC , which is the steady state value of the aggregate load as shown in Proposition 1, for arbitrary initial condition $X^{(0)}$ and therefore the aggregate load is globally stable.

Now we show that the convergence of the aggregate load implies the convergence of the individual rates. With $a^{(k)} := 1 - \kappa p^{(k)}$ and $b_i := \kappa w_i, \forall i \in \mathcal{S}$, we can rewrite the sender equation as

$$x_i^{(k)} = a^{(k-1)} \cdot x_i^{(k-1)} + b_i.$$

Applying Lemma 1, we obtain an explicit form for $x_i^{(k)}$:

$$x_i^{(k)} = x_i^{(0)} \cdot \prod_{m=0}^{k-1} a^{(m)} + b_i \cdot \sum_{m=0}^{k-1} \prod_{p=m+1}^{k-1} a^{(p)}. \quad (7)$$

Since we know that $X^{(k)} \rightarrow \mu C$, we conclude that there exists a $k_0 \in \mathbb{N}$, such that all $a^{(k)}$ with $k \geq k_0$ lie in the interval $(0, 1)$:

$$\begin{aligned} X^{(k)} &\rightarrow \mu C \\ \Rightarrow p^{(k)} &\rightarrow \frac{\mu - \nu}{\mu} \in (0, 1) \\ \Rightarrow 0 < \epsilon_1 \leq p^{(k)} < 1, &\quad \forall k > k_0 \\ \Rightarrow 0 < a^{(k)} \leq \epsilon_2 < 1, &\quad \forall k > k_0. \end{aligned}$$

From this we derive two conclusions. First:

$$x_i^{(0)} \cdot \prod_{m=0}^{k-1} a^{(m)} \xrightarrow[k \rightarrow \infty]{} 0. \quad (8)$$

And second:

$$b_i \cdot \sum_{m=0}^{k-1} \prod_{p=m+1}^{k-1} a^{(p)} \quad (9)$$

is strictly increasing for $k \geq k_0$. Since we know that $X^{(k)}$ converges, we also know that $x_i^{(k)}$ is bounded $\forall i \in \mathcal{S}$. From this fact, using (7) and (8), we know that sequence (9) is bounded too. Since (9) is strictly increasing and bounded, it converges, and using (7) we obtain convergence of $x_i^{(k)}$.

Now that we know that $x_i^{(k)}$ converges, we can take the limit on both sides of the sender equation and obtain $\lim_{k \rightarrow \infty} x_i^{(k)} = \frac{\mu - \nu}{\mu} w_i$, which is the steady state value of x_i as shown in Proposition 1. ■

3) *Multi-Link Stability*: As argued in [1], proofs of multi-link stability under general conditions are very challenging. However, with the assumption of a fixed consistent bottleneck assignment (see [1] for a definition), Zhang, Kang and Loguinov [Theorem 3 in [1]] showed that for certain max-min systems including EMKC local stability in single-link networks implies local stability in the multi-link case. A consistent bottleneck assignment with EMKC and xMKC means that each source receives feedback from the router on its path that has the highest congestion measure $p_l(t)$. Since the proof of this theorem does not rely on the assumption of $w_i = w_j, \forall i, j \in \mathcal{S}$, single-link results of Theorem 1 imply multi-link stability of xMKC with a fixed consistent bottleneck assignment in the presence of heterogeneous time-varying feedback delays.

IV. UTILITY MAX-MIN FAIRNESS WITH UMKC

With xMKC, a target steady state rate allocation is determined by an arbitrary but fixed parameter vector w , see Proposition 1. In a distributed setting, this parameter vector w is not known a priori.

In this section, we introduce *available utility* u as a second feedback variable, which allows the senders to dynamically adapt w_i in a distributed way so as to achieve a utility max-min fair rate allocation in the steady state. We propose a strategy for a flow to choose its feedback router according to both feedback variables and call the extended algorithm UMKC. We prove that UMKC's steady state is utility max-min fair in arbitrary networks. Based on the results for xMKC, we establish stability results for UMKC in the presence of time-varying feedback delays. Note that our analysis does not rely on strict concavity of utility functions.

A. Utility Max-Min Fair Steady State

Cao and Zegura [29] established the following necessary and sufficient condition for a rate allocation to be utility max-min fair:

Proposition 3 (from Cao and Zegura [29]). *A feasible rate allocation $x = (x_i, i \in \mathcal{S})$ is utility max-min fair if and only if each flow has a utility bottleneck link with respect to x .*

Here, utility bottleneck is defined as follows

Definition 1 (from Cao and Zegura [29]). *Given a feasible bandwidth allocation $x = (x_i, i \in \mathcal{S})$, we say that for a flow $i \in \mathcal{S}$ link $l \in \mathcal{L}(i)$ is a utility bottleneck if and only if all of l 's target capacity is allocated and $U_i(x_i) \geq U_j(x_j), \forall j \in \mathcal{S}(l), j \neq i$.*

In [29], the authors also showed that for a given set of utility functions and a given network there always exists a utility max-min fair rate allocation.

To achieve utility max-min fairness in the steady state, we use the following parameter vector:

$$w_i = (1 - \nu)U_i^{-1}(u), \quad i \in \mathcal{S}. \quad (10)$$

Here, u denotes the available utility of the path, which is part of the feedback provided by the network. We obtain the following rate equation:

$$x_i(t) = x_i(t - d_i) + \kappa [(1 - \nu)U_i^{-1}(u_{l(i)}) - \eta_i(t) \cdot x_i(t - d_i)]. \quad (11)$$

Here, $\eta_i(t) := p_{l(i)}(t - d_{i,l(i)}^-)$ and $l(i)$ is the link on the path of the flow i that provides the feedback.

With UMKC, the feedback router for flow i is determined according to the following scheme:

Scheme 1 Selection of the feedback router with UMKC

1. If on the path of flow i there are overloaded links, the router with the highest congestion measure provides the feedback: $l(i) = \operatorname{argmax}_{l \in \mathcal{L}(i)} (p_l(t - d_{i,l}^-))$, $\forall i \in \mathcal{S}$.
 2. Otherwise, the feedback is provided by the router with the smallest available utility u_l on the path: $l(i) = \operatorname{argmin}_{l \in \mathcal{L}(i)} (u_l)$, $\forall i \in \mathcal{S}$.
 3. If there are multiple routers with the smallest available utility on the path, choose the one with the highest router ID. (This rule is needed for consistency and is used in the proof of Lemma 2.)
-

To implement this scheme, each router is required to examine the congestion measure and the available utility that is already contained in the header of arriving packets and replace them with its own values if needed.

Following [1], we call a flow i that receives its feedback from the link $l \in \mathcal{L}(i)$ *responsive* with respect to l and *unresponsive* with respect to all other links $l' \in \mathcal{L} \setminus \{l\}$.

In the following Proposition we show that using source equation (11) and selection Scheme 1 and provided that all routers that have responsive flows are fully utilized, the resulting rate allocation is utility max-min fair.

Proposition 4. *In an arbitrary network, if the rate allocation is feasible and all routers that have responsive flows are fully utilized then the steady state of the system (2), (11), with feedback router selection Scheme 1, is utility max-min fair.*

Proof: A necessary and sufficient condition for (11) to be in its steady state is $\hat{x}_i = \frac{(1-\nu)U_i^{-1}(u_{l(i)})}{\hat{p}_{l(i)}}$, where $\hat{p}_{l(i)}$ is the steady state congestion measure of link $l(i)$. Since we assumed that all routers that have responsive flows are fully utilized, we obtain from (2): $\hat{p}_{l(i)} = 1 - \nu$, which implies that:

$$\hat{x}_i = U_i^{-1}(u_{l(i)}). \quad (12)$$

Since we also assumed that the rate allocation is feasible, each flow receives its feedback from the router with the smallest available utility on the path according to step 2 in Scheme 1. To complete the proof, we need to show that each flow has a utility bottleneck, so that the conditions of Proposition 3 are satisfied. Since we assumed that each router that has responsive flows is fully utilized, it remains to show that for each flow i : $U_j(\hat{x}_j) \leq U_i(\hat{x}_i), \forall j \in \mathcal{S}(l(i))$. With (12),

this translates to $u_{l(j)} \leq u_{l(i)}$, $\forall j \in \mathcal{S}(l(i))$. Assume on the contrary that there exists a $j \in \mathcal{S}(l(i))$ such that $u_{l(j)} > u_{l(i)}$. We immediately obtain a contradiction to step 2 in Scheme 1, since both $l(i)$ and $l(j)$ are on the path of flow j and flow j receives its feedback from router $l(j)$ though the router $l(i)$ has lower available utility. ■

Now we know that to establish a utility max-min fair rate allocation we must ensure that no router is overloaded and each router that has responsive flows is fully utilized. The only way for the routers to influence their aggregate load is through providing appropriate feedback u . Since the routers are not able to directly calculate the desired u as it requires the knowledge of the utility functions of all flows they are traversed by, they have to update their available utility iteratively, generating a sequence $u_l^{(\tau)}$ that converges to the desired value, which we denote by \hat{u}_l . The existence of this value is ensured by the result from [29] that states that for each set of utility functions and for each network there is a utility max-min fair rate allocation. After each update of u , the routers have to wait until the sending rates reach their corresponding steady state values to avoid oscillations of the available utility and inconsistent selection of feedback routers by the flows. Then, they can recalculate u based on the measured aggregate load $\hat{X}(u)$. We thus obtain a sequence of steady states. To simplify the notation, we call the steady states that correspond to $u \neq \hat{u}$ transient steady states and the utility max-min fair steady state with $u = \hat{u}$ terminal steady state.

For the calculation of the available utility we use the secant method. Writing $\hat{X}_l^{(\tau)}$ for $\hat{X}_l(u_l^{(\tau)})$, it has the shape:

$$u_l^{(\tau)} = u_l^{(\tau-1)} - \left[\hat{X}_l^{(\tau-1)} - C_l \right] \frac{u_l^{(\tau-1)} - u_l^{(\tau-2)}}{\hat{X}_l^{(\tau-1)} - \hat{X}_l^{(\tau-2)}}. \quad (13)$$

Since this method works only for routers that have responsive flows (because otherwise \hat{X}_l is not a function of u_l), we assume in the following that all routers have responsive flows. Note that in the case of linear utility functions, $\hat{X}(u)$ is also a linear function of u and the secant method converges in only one step. This implies that in a single-link network with linear utility functions, starting with an arbitrary $u > 0$, the router is able to calculate \hat{u} after two transient steady states that provide the values $\hat{X}^{(0)}$ and $\hat{X}^{(1)}$. The third steady state will be utility max-min fair with $\hat{X}^{(3)} = C$.

Now we are able to prove the following Theorem:

Theorem 3. *The terminal steady state of UMKC as described by equations (2), (11), and (13), and with the feedback router selection Scheme 1, is utility max-min fair in an arbitrary multi-link network, where all routers have responsive flows.*

Proof: Consider equation (13). It is in its steady state if and only if $\hat{X}_l = C_l$. Since we assume that all routers have responsive flows, the resulting rate allocation is feasible and the conditions of Proposition 4 are fulfilled. ■

In the following, we present a result showing that for utility functions of the form $U_i(x) = x$, $\forall i \in \mathcal{S}$, UMKC leads to

bandwidth max-min fairness.

Corollary 1. *The terminal steady state of UMKC as described by equations (2), (11), and (13), with the feedback router selection Scheme 1, and with utility functions satisfying $U_i(x) = x$, $\forall i \in \mathcal{S}$, is max-min fair in an arbitrary multi-link network, where all routers have responsive flows.*

B. Single-Link Stability of UMKC

In this section, we use the results established for xMKC to derive a sufficient condition for the transient and the terminal steady state of UMKC to be stable in single-link networks. First, we establish the equilibrium structure of UMKC in a single-link network:

Proposition 5. *In a single-link network, steady state of the system (2), (11) with fixed $\nu \in (0, 1)$, $\kappa \in (0, 1)$ and $u > 0$ has the following shape:*

$$\hat{x}_i(u) = \frac{U_i^{-1}(u)}{\sum U_i^{-1}(u)} \hat{X}(u) \quad (14)$$

$$\hat{X}(u) = (1 - \nu) \sum U_i^{-1}(u) + \tilde{C} \quad (15)$$

Corollary 2 (to Theorem 1). *In a single-link network with heterogeneous time-varying feedback delays, the system (2), (11) with $\nu \in (0, 1)$ and $\kappa \in (0, 1)$ is locally asymptotically stable in its transient steady state with $u > 0$ if $\hat{X}(u) \geq 2\nu C$. For the terminal steady state, this condition translates to $\nu \leq 0.5$.*

Note that as long $u \neq \hat{u}$, $\hat{X}(u)$ can become arbitrarily small. Therefore, it is not possible to choose ν such that transient steady states for all $u > 0$ are locally asymptotically stable in the presence of heterogeneous time-varying feedback delays. For example, choosing $\nu = 0.5$ guarantees stability only for the terminal steady state with $\hat{X}(u) = C$. Choosing $\nu = 0.4$ guarantees stability for all transient steady states with total load $\hat{X}(u) \geq 0.8C$. However, this established stability condition is only proved to be sufficient. Its necessity is still an open issue. In fact, simulations show that UMKC is still stable even if this condition is violated.

Similarly, global asymptotic stability of UMKC with a fixed $u > 0$ follows directly from the corresponding statements for xMKC.

Corollary 3 (to Theorem 2). *In a single-link network with homogeneous time-varying feedback delays, system (2), (11) with $\nu \in (0, 1)$ and $\kappa \in (0, 1)$ is globally asymptotically stable for an arbitrary fixed $u > 0$.*

This result holds for a fixed $u > 0$, which corresponds to a transient steady state if $u \neq \hat{u}$ and to the terminal steady state if $u = \hat{u}$. It implies that UMKC globally converges towards the terminal steady state if and only if the sequence of transient steady states converges. In the following, we present a special class of utility functions, which guarantees the convergence of the sequence of transient steady states.

Theorem 4. *In a single-link network with homogeneous time-varying feedback delays and linear utility functions, UMKC is globally asymptotically stable.*

Proof: Global stability of UMKC requires two conditions: (i) global stability of the sequence of $u^{(\tau)}$ of available utilities and (ii) global stability of each transient steady state (corresponding to a fixed $u^{(\tau)}$). The second condition directly follows from Corollary 3. The first condition follows from global convergence of the secant method for linear functions. ■

As a special case of the above theorem, for utility functions satisfying $U_i(x) = x, \forall i \in \mathcal{S}$, UMKC is globally stable. This case is of independent interest as it results in a max-min fair rate allocation (see Corollary 1).

C. Multi-Link Stability of UMKC

In this section we show that UMKC is locally asymptotically stable in its terminal steady state in multi-link networks with heterogeneous time-varying feedback delays if each sender is able to correctly identify its utility bottleneck link. We use a technique presented by Zhang, Kang and Loguinov [1]. Further investigations of UMKC's behavior in multi-link networks are made by means of simulations.

Following [1], we say that the steady state rate allocation of flow j depends on the steady state rate allocation of flow i or simply flow j depends on flow i if and only if $l(j) \in \mathcal{L}(i) \cap \mathcal{L}(j)$ and $l(i) \neq l(j)$. Here, $l(i) \in \mathcal{L}(i)$ denotes again the feedback router of flow $i, \forall i \in \mathcal{S}$. We denote the fact that flow j depends on flow i by $i \rightarrow j$. This dependency relation allows us to construct a dependency graph for a given assignment of feedback routers. Analogous to [1], we prove the following proposition for UMKC

Lemma 2. *With UMKC, if each sender receives feedback from its utility bottleneck link, then the resulting dependency graph is acyclic.*

Proof: From its definition we know that a utility bottleneck router is fully utilized. As shown in the proof of Proposition 4, the received utility of a flow equals the available utility of its feedback router if the feedback router is fully utilized: $U_i(x_i) = u_{l(i)}$. Since for the utility bottleneck of flow i we know: $U_i(x_i) \geq U_j(x_j), \forall j \in \mathcal{L}(l(i))$, we conclude

$$u_{l(i)} \geq u_{l(j)}, \quad \forall j \in \mathcal{L}(l(i)). \quad (16)$$

Now assume that there is a cycle in the dependency graph: $i_1 \rightarrow \dots \rightarrow i_n \rightarrow i_1$. With (16), step 2 in the feedback router selection Scheme 1 and using the definition of the introduced dependency relation, we obtain: $u_{l(i_1)} \leq \dots \leq u_{l(i_n)} \leq u_{l(i_1)}$. Now, consider two cases. First, one of the inequality signs is strict. Then, we obtain the contradiction $u_{l(i_1)} < u_{l(i_1)}$ and the proof is complete. Second, all available utilities are equal. In this case, step 3 in the feedback router selection Scheme 1 demands that the router with the highest router ID provides the feedback. We obtain: $l(i_1) < \dots < l(i_n) < l(i_1)$, thus obtaining the contradiction $l(i_1) < l(i_1)$. ■

Now, following [1], we call an assignment of feedback routers *consistent* if and only if it results in an acyclic dependency graph. To prove the following stability result, we use a result from [1] showing that a network with a consistent assignment of feedback links contains at least one link that has no unresponsive flows.

Theorem 5. *In a multi-link network and provided that all routers have responsive flows, each flow receives feedback from its utility bottleneck link, and $\nu \leq 0.5, u = \hat{u}$ are fixed, UMKC is locally asymptotically stable in its terminal steady state in the presence of heterogeneous time-varying feedback delays.*

Proof: From Lemma 2, we know that if each flow receives its feedback from its utility bottleneck then the resulting dependency graph is acyclic and therefore the assignment of the feedback routers is consistent. From Lemma 3 in [1], we know that with a consistent assignment of feedback routers, there is at least one router l_1 that has no unresponsive flows. Therefore, we can consider this router and its responsive flows as a single-link network, since both are independent from the other flows. For single-link networks, we know from Corollary 2 that system (2), (11), which describes UMKC with fixed $u > 0$, is locally asymptotically stable in the terminal steady state, that is with $u = \hat{u}$, if $\nu \leq 0.5$. After the flows that are responsive with respect to l_1 have reached their steady state, we can remove them from the dependency graph. Again, the new dependency graph contains at least one router l_2 that has no unresponsive flows except for the flows of the router l_1 that we now can assume as constant-rate. Repeating this argumentation chain for all routers in the network, we obtain asymptotic local stability of the entire system. ■

This Theorem proves local asymptotic stability of UMKC's terminal steady state in multi-link networks with a stationary assignment of feedback routers. It does not address the issue of oscillations of the assignment of feedback routers. During such oscillations, there is a possibility of having a directed cycle in the dependency graph that persists over time so that the proof is not valid for that case. Feedback might oscillate in the presence of heterogeneous feedback delays when senders update their rates based on the information that was valid at different time instants. We leave it for future work to prove the convergence of UMKC's assignment of feedback routers.

V. SIMULATION RESULTS

We extended the NS-2 framework [2] to simulate UMKC in multi-link topologies. In this section we present the results of the simulations.

A. Single-Link Simulation

Figure 1 shows the results of a single-link simulation with two flows, one with a strict concave utility function $U_1(x) = k\sqrt{x}$ and the other with a multimedia utility function $U_2(x)$ as in Figure 1a. The multimedia application is assumed to use a codec that operates optimally at 0.5 Mbps and quickly degrades in performance if the available bandwidth decreases




Fig. 1. Results of the single-link simulation: (a) Used utility functions. (b) Sending rates and their sum as percentage of the target load.

(e.g., a video stream). The multimedia utility function is constructed from linear and tanh sections in a way that it has a specified minimum and maximum slope. The target capacity of the link is 0.6 Mbps, which is 90% of the physical capacity to prevent large queues.

With the bandwidth fair resource allocation, both applications would receive 0.3 Mbps. This is not enough for the multimedia application to start transmission, therefore its received utility would be near zero. In contrast to this, utility max-min fair bandwidth allocation gives 0.18 Mbps to the elastic application and 0.42 Mbps to the multimedia application, which allows the latter to start sending though with a reduced performance. Figure 1b shows the sending rates of the flows as a percentage of the target capacity. Flow 1 starts at $t = 0$ sec and flow 2 starts at $t = 5$ sec. Both sending rates converge to their predicted utility max-min fair values of approximately 30% and 70% of the target load.

B. Multi-Link Simulation

We evaluated UMKC in a multi-link topology given in Figure 2a with heterogeneous feedback delays and randomly arriving and departing flows to investigate its stability, fairness, link utilization and queue length. For this simulation, we used three utility functions, a concave utility function $U_c(x) = k_c\sqrt{x}$, a linear one $U_l(x) = k_lx$, and a multimedia utility function of a video stream $U_m(x)$ with a base layer at 0.5 Mbps and two enhancement layers at 1 Mbps and 1.5 Mbps as in Figure 2b.

The target load of Link 1 is 20 Mbps, the one of Link 2 is 25 Mbps. For both links, target loads are 90% of the physical capacities to prevent building of large queues. The delay of Link 1 is 1 ms, the one of Link 2 is 100 ms. Additionally, each flow has a random round trip delay between 1 ms and 100 ms.

At $t = 0$, 11 senders are attached to the SN1. 10 start to send to DN2 and 1 to DN1. At $t = 30$ sec, further 10 senders join the network at SN2 and send to DN2. At $t = 50$ sec, senders start to join the network at random time instants chosen from the interval $[50, 100]$ sec. In total, 20 senders join the network during this time, 10 at SN1 sending to DN2 and 10 at SN2 sending to DN2. At $t = 100$ sec, senders start to leave network at random time instants chosen from the interval $[100, 150]$ sec. 20 senders leave the network during this time, 10 from SN1




Fig. 2. (a) Three core nodes (CN) are interconnected by two links with target capacities C_1 and C_2 . There are 3 sets of flows that interconnect a source node (SN) and a destination node (DN): $SN1 \rightarrow DN1$, $SN1 \rightarrow DN2$ and $SN2 \rightarrow DN2$. Each flow is randomly assigned a round trip delay between 1 ms and 200 ms, so that we obtain a heterogeneously delayed network. (b) Video stream utility function used for the multi-link simulation. It is constructed from linear and tanh sections and its slope lies in the interval $[0.1, 50]$.

and 10 from SN2. These senders are chosen randomly from the whole set of the senders.

50% of the flows have the strict concave utility function $U_c(x)$, 40% the linear utility function $U_l(x)$ and 10% the multimedia utility function $U_m(x)$.

Figure 3a shows the fairness index adopted from Jain, Chiu and Hawe [38]: $\phi(x) = \frac{[\sum_{i \in S} U_i(x_i)]^2}{|S| \sum_{i \in S} U_i(x_i)^2}$. Here, $U_i(x_i)$ is the received utility of flow i and $|S|$ is the total number of flows in the network. The range of values of $\phi(x)$ is $(\frac{1}{|S|}, 1]$, where 1 means that all flows receive the same utility and $\frac{1}{|S|}$ means maximum unfairness. It is remarkable that the fairness level is above 0.9 even during the period with randomly arriving and departing flows. At $t = 30$ sec, 10 flows enter the network simultaneously, which temporarily degrades the fairness. Also at other time instants, when multiple flows enter the network in quick succession, the fairness level is temporarily decreased.

VI. CONCLUSION

In this paper, we addressed two important goals in designing congestion control mechanisms: delay-independent stability in presence of heterogeneous time-varying feedback delays and joint congestion control of flows with utility functions that are not necessarily concave. We achieved the first goal by proposing a flexible framework called xMKC, which can be used for designing congestion control algorithms that achieve delay-independent stability and allocate resources fairly. We achieved the second goal by proposing the congestion control algorithm UMKC, which is based on the previously established framework and leads to a utility max-min fair rate allocation in arbitrary networks. Based on the results for xMKC, we established stability results for UMKC in the presence of time-varying feedback delays. We proved asymptotic global stability for time-varying delays in single-link networks for a max-min fair congestion controller, which is an important special case of our framework.

An open issue that we leave for future work is the investigation of stability characteristics without the assumption of a fixed and consistent assignment of the feedback routers (see also [39]). Another issue is a unified normalization framework

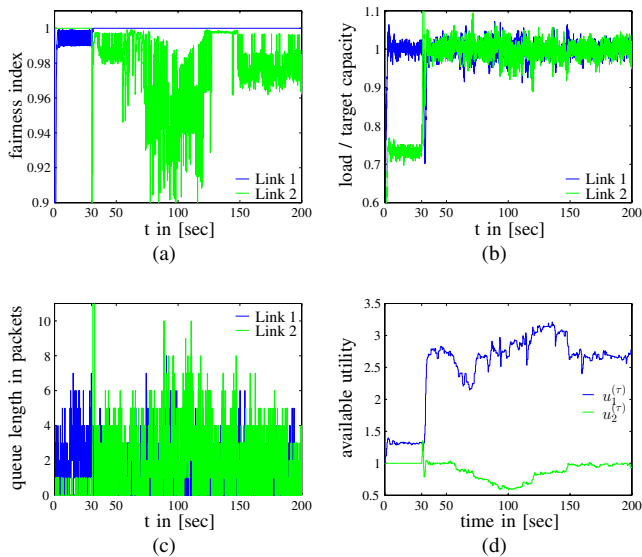


Fig. 3. Results of the multi-link simulation: (a) Fairness index adopted from Jain, Chiu and Hawe [38]. (b) The quotient of the aggregate load and the target load. (c) Queue lengths. (d) The available utility $u_i^{(\tau)}$.

for utility functions (see also [40]). Finally, we see potential in the systematic investigation of discrete congestion control algorithms satisfying the stability condition $\rho(|\mathcal{J}|) < 1$.

ACKNOWLEDGMENT

The authors would like to thank Michael Karow, Tobias Poschwatta and Yueping Zhang for helpful suggestions and discussions.

REFERENCES

[1] Y. Zhang, S.-R. Kang, and D. Loguinov, "Delay-independent stability and performance of distributed congestion control," *IEEE/ACM Trans. on Networking*, vol. 15, 2007.

[2] S. McCanne and S. Floyd, "The network simulator - ns-2," <http://www.isi.edu/nsnam/ns>.

[3] V. Jacobson and M. Karels, "Congestion avoidance and control," *ACM SIGCOMM*, vol. 18, pp. 273–288, 1988.

[4] V. Jacobson, "Modified tcp congestion avoidance algorithm," April 1990, message to end2end-interest mailing list.

[5] S. Floyd and T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm," RFC 2582.

[6] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP Selective Acknowledgment Options," RFC 2018.

[7] L. S. Brakmo and L. L. Peterson, "Tcp vegas: End to end congestion avoidance on a global internet," *IEEE JSAC*, vol. 13, pp. 1465–1480, 1995.

[8] S. H. Low and D. E. Lapsley, "Optimization flow control i: Basic algorithm and convergence," *IEEE/ACM Trans. on Networking*, vol. 7, pp. 861–874, 1999.

[9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness, and Stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–52, 1998.

[10] S. Low, F. Paganini, and J. C. Doyle, "Internet Congestion Control," *IEEE Control Systems Magazine*, vol. 22, 2002.

[11] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. on Networking*, vol. 8, pp. 556–567, 2000.

[12] B. Wyrowski, L. L. H. Andrew, and M. Zukerman, "Maxnet: A congestion control architecture for scalable networks," *IEEE Communications Letters*, vol. 7, pp. 511–513, 2003.

[13] S. Deb and R. Srikant, "Global stability of congestion controllers for the internet," *IEEE Trans. on Automatic Control*, vol. 48, pp. 1055–1060, 2003.

[14] R. Johari and D. K. H. Tan, "End-to-end congestion control for the internet: Delays and stability," *IEEE/ACM Trans. on Networking*, vol. 9, pp. 818–832, 2001.

[15] K. Kar, S. Sarkar, and L. Tassiulas, "A simple rate control algorithm for maximizing total user utility," in *IEEE INFOCOM*, 2001.

[16] D. Katabi, M. Handley, and C. Rohrs, "Congestion control for high bandwidth delay product networks," in *ACM SIGCOMM*, 2002.

[17] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue (avq) algorithm for active queue management," in *ACM SIGCOMM*, 2001.

[18] —, "A time-scale decomposition approach to adaptive explicit congestion notification (ecn) marking," *IEEE Trans. on Automatic Control*, vol. 47, pp. 882–894, 2002.

[19] —, "End-to-end congestion control schemes: Utility functions, random losses and ecn marks," *IEEE/ACM Trans. on Networking*, vol. 11, pp. 689–702, 2003.

[20] —, "Stable, scalable, fair congestion control and aqm schemes that achieve high utilization in the internet," *IEEE Trans. on Automatic Control*, vol. 48, pp. 2024–2029, 2003.

[21] L. Massoulié, "Stability of distributed congestion control with heterogeneous feedback delays," *IEEE/ACM Trans. on Networking*, vol. 47, pp. 895–902, 2002.

[22] G. Vinnicombe, "On the stability of end-to-end congestion control for the internet," *Cambridge University, Technical Report CUED/FINFENG/TR.398*, Dec. 2000.

[23] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhäuser Boston, 2003.

[24] Y. Zhang and D. Loguinov, "On delay-independent diagonal stability of max-min congestion control," in *IEEE CDC*, 2006.

[25] S. Shenker, "Fundamental Design Issues for the Future Internet," *IEEE JSAC*, vol. 13, pp. 1176–88, 1995.

[26] L. Breslau and S. Shenker, "Best-effort versus reservations: a simple comparative analysis," in *ACM SIGCOMM*, 1998.

[27] K. Lee, "Adaptive network support for mobile multimedia," in *ACM MOBICOM*, 1995.

[28] D. Reininger, "Dynamic quality-of-service framework for video in broadband networks," Ph.D. dissertation, Rutgers University, New Jersey, 1998.

[29] Z. Cao and E. Zegura, "Utility max-min: An application-oriented bandwidth allocation scheme," in *IEEE INFOCOM*, 1999, pp. 793–801.

[30] T. Harks, "Utility Proportional Fair Resource Allocation: An Optimization Oriented Approach," in *Proc. of QoS in Multiservice IP Networks*, 2005.

[31] J. W. Lee, R. R. Mazumdar, and N. B. Shroff, "Non-convex optimization and rate control for multi-class services in the internet," *IEEE/ACM Trans. on Networking*, vol. 13, pp. 1063–1092, 2005.

[32] M. Chiang, S. Zhang, and P. Hande, "Distributed rate allocation for inelastic flows: optimization frameworks, optimality conditions, and optimal algorithms," in *IEEE INFOCOM*, 2005.

[33] S. Sarkar and L. Tassiulas, "Fair allocation of utilities in multirate multicast networks: A framework for unifying diverse fairness objectives," *IEEE Trans. on Automatic Control*, vol. 47, pp. 931–944, 2002.

[34] H.-W. Lee and S. Chong, "A Distributed Utility Max-Min Flow Control Algorithm," *Computer Networks*, vol. 50, pp. 1816–1830, 2006.

[35] J. Cho and S. Chong, "Utility max-min flow control using slope-restricted utility functions," *IEEE Trans. on Communications*, vol. 55, pp. 963–972, 2007.

[36] E. Kaszkurewicz and A. Bhaya, "A delay-independent robust stability condition for linear discrete-time systems," *IEEE CDC*, vol. 4, pp. 3459–3460, 1993.

[37] Y. Zhang and D. Loguinov, "Local and global stability of symmetric heterogeneously-delayed control systems," in *IEEE CDC*, 2004.

[38] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," Digital Equipment Corporation, DEC Research Report TR-301, 1984.

[39] Y. Zhang, D. Leonard, and D. Loguinov, "Jetmax: Scalable max-min congestion control for high-speed heterogeneous networks," in *IEEE INFOCOM*, 2006.

[40] T. Harks and T. Poschwatta, "Priority pricing in utility fair networks," in *IEEE ICNP*, 2005.