Lyapunov-Based Stability Analysis for REM Congestion Control

Orhan Ç. IMER, Tamer BAŞAR

***Abstract*— This paper investigates convergence properties of basic REM flow control algorithm via Lyapunov functions. The decentralized algorithm REM consists of a link algorithm that updates a congestion measure, also called "price", based on the excess capacity and backlog at that link, and a source algorithm that adapts the source rate to congestion in its path. At the equilibrium of the algorithm, links are fully utilized, and all buffers are cleared. Convergence of the algorithm is established for single and two-link cases using a Lyapunov argument. Extension to the general multi-link model is discussed as well.**

I. Introduction

In general, network flow control concerns adjustment of individual transmission rates of a number of sources over a set of links, subject to link capacity constraints. The main purpose of flow control is to fully utilize all the links in the network, while at the same time achieving some sort of fairness among the sources. There are several ways of defining fairness in a network, each one leading to a different allocation of link capacities, see [1].

In optimization based flow control, allocation of link bandwidths is achieved through a utility maximization problem, in which each source has a utility function reflecting its valuation of transmitting at a certain rate. The aggregate source utility is then maximized over transmission rates of all sources subject to the capacity constraints. The significance of the utility maximization model is that TCP Reno [2] and Vegas [3], as well as several other Internet congestion control algorithms can be interpreted within this model by choosing appropriate utility functions. This approach has been extensively studied in several papers, such as [4], [5], [6], [7], [8]. Solving the utility maximization problem directly requires coordination among possibly all sources, and hence is impractical for the networks. However, there exist decentralized solutions which only use partial information from the network. One way of obtaining a decentralized solution is to look at the dual of the maximization problem. In [6], it has been shown that this leads to a decentralized algorithm that consists of a link algorithm that updates a congestion measure, also called "price", based on the excess capacity at that link, and a source algorithm that adapts the source rate to congestion in its path. This algorithm, which is shown to converge, has a drawback in that the backlog can be quite large in the equilibrium. To remedy this problem, in [9], REM (Random Exponential Marking) algorithm has been introduced, which ensures that the buffer is cleared when the equilibrium is reached. The convergence of this more appealing algorithm for a multi-link network has not been established yet. In [10], the continuous time version of REM has been shown to be globally stable, but for the original discrete-time case, a proof of stability is available only for the single-link case, see [11]. The stability analysis of [11] is based on an invariance argument which does not generalize to multiple links. In this paper, we prove a similar stability result for a single-link network using a Lyapunov argument. The advantage of using a Lyapunov function is that it can be generalized to a multi-link network. However, due to cumbersome algebra for a network of size more than two-links, in this paper, we only present the stability proof for a two-link network. Extension to the multi-link case is also discussed.

The rest of the paper is organized as follows. In Section 2, the network model is introduced. The stability analysis is carried out in Section 3. In Section 4 we discuss modifying the REM algorithm to make it more suitable for Lyapunov-type stability analysis. The paper ends with the concluding remarks of Section 4, in which we also discuss some future research directions.

II. The Network Model

Consider a network $\mathcal{N}$ that consists of a set $\mathcal{L} = \{1, \ldots, L\}$ of links of capacity c_l, $l \in \mathcal{L}$. The network is shared by a set $\mathcal{S} = \{1, \ldots, S\}$ of sources. Source s transmits at rate x_s using a set $\mathcal{L}_s \subseteq \mathcal{L}$ of links. The routing matrix R, of dimension $L \times S$, is defined by $R_{ls} = \mathcal{X}_{s \in \mathcal{S}_l}$, where $\mathcal{S}_l$ is the set of sources using link l, and $\mathcal{X}_A$ denotes the indicator function of the set A.

The rate x_s satisfies $m_s \leq x_s \leq M_s$, where $m_s \geq 0$ and $M_s < \infty$ are the minimum and maximum transmission rates, respectively. When transmitting at rate x_s, source s attains a utility of $U_s(x_s)$. It is assumed that the utility functions U_s are strictly concave increasing and twice continuously differentiable. Associated with each link l there is a buffer with occupancy b_l. The amount of backlog at link l satisfies $0 \leq b_l \leq B_l$, where $B_l < \infty$ is the maximum buffer occupancy.

Department of Electrical and Computer Engineering and Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801-2307, USA. Email: imer@control.csl.uiuc.edu.

Department of Electrical and Computer Engineering and Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801-2307, USA. Tel: (217) 333-3607; Fax: (217) 244-1653; Email: tbasar@control.csl.uiuc.edu Research supported by the NSF grant NSF CCR 00-85917

We assume an underlying discrete time structure with no delay and synchronous updates. Let t denote the discrete time unit, and assume that fluid approximation for queue lengths holds. Then, the buffer occupancy $b_l^{(t)}$ at link l at time t evolves according to

$$b_l^{(t+1)} = \left[b_l^{(t)} + \sum_{s \in \mathcal{S}_l} x_s^{(t)} - c_l \right]_0^{B_l} \tag{1}$$

where the notation $[x]_{x_{\min}}^{x_{\max}}$ is shorthand for

$$[x]_{x_{\min}}^{x_{\max}} = \begin{cases} x_{\min} & x \le x_{min} \\ x & x_{\min} \le x \le x_{\max} \\ x_{\max} & x \ge x_{\max} \end{cases}$$

Our objective (the primal problem) is to choose the source rates x_s so as to:

$$\max_{m_s \le x_s \le M_s} \sum_{s \in \mathcal{S}} U_s(x_s) \tag{2}$$

subject to capacity constraints:

$$\sum_{s \in \mathcal{S}_l} x_s \le c_l \tag{3}$$

This flow control problem is first posed in [4], and solved in [5] using techniques of constrained optimization. The problem can also be solved using a penalty function approach, see [7]. Note that a unique maximizer exists, called the primal optimal solution, since the objective function is strictly concave, and the feasible solution set is compact. Even though the objective function is separable in x_s, the source rates are coupled by the constraint (3). Thus, solving the primal problem directly requires coordination among possibly all sources and is impractical for networks. Looking at the dual of the optimization problem (2)-(3), one can obtain a gradient projection algorithm in the dual variables that lends itself into a decentralized implementation [6]. The Lagrangian dual of this problem is

$$\min_{p_l \ge 0} \max_{m_s \le x_s \le M_s} \sum_{s \in \mathcal{S}} U_s(x_s) - \sum_{l \in \mathcal{L}} p_l \Big(\sum_{s \in \mathcal{S}_l} x_s - c_l \Big) \tag{4}$$

where associated with each link there is a dual variable p_l, termed as "price" of link l. The solution of this dual optimization problem can be obtained using a gradient projection algorithm. The basic algorithm given in [6] works as follows. At times $t = 1, 2, \ldots$, link l receives rates from all sources $s \in \mathcal{S}_l$ that go through it, and it calculates the aggregate link rate $y_l^{(t)} := \sum_{s \in \mathcal{S}_l} x_s^{(t)}$. With this information, it computes the so-called link price $p_l^{(t)}$, which is updated according to

$$p_l^{(t+1)} = [p_l^{(t)} + \gamma (\sum_{s \in \mathcal{S}_l} x_s^{(t)} - c_l)]_0^{\infty} \tag{5}$$

where $\gamma > 0$ is a step size. Note that this update can be interpreted as a gradient step in the minimization of (4). This price is communicated to all sources $s \in \mathcal{S}_l$ that use link l. Source s upon receiving the aggregate price $q_s^{(t)} := \sum_{l \in \mathcal{L}_s} p_l^{(t)}$ of all links in its path, chooses a new transmission rate $x_s^{(t)}$ by solving the maximization problem

$$\max_{m_s \le x_s \le M_s} U_s(x_s) - x_s q_s \tag{6}$$

whose solution is trivially given by $x_s^{(t)} = [U_s'^{-1}(q_s^{(t)})]_{m_s}^{M_s}$. Here $U_s'^{-1}$ denote the inverse functions of the marginal utility; they exist and are strictly decreasing because U_s are strictly concave increasing.

In [6], it is shown that this distributed algorithm converges to the solution of the primal optimization problem (2)-(3) if the curvatures of U_s are bounded away from zero, i.e. $-U_s''(x_s) \ge -1/\rho_s > 0$, $m_s \le x_s \le M_s$, and if the step size γ satisfies $0 < \gamma < 2/\rho \bar{L} \bar{S}$, where $\rho = \max_{s \in \mathcal{S}} \rho_s$, $\bar{L} = \max_{s \in \mathcal{S}} |\mathcal{L}_s|$, and $\bar{S} = \max_{l \in \mathcal{L}} |\mathcal{S}_l|$.

Note however that here prices integrate excess capacity, which is exactly what the backlog does:

$$b_l^{(t+1)} = \left[b_l^{(t)} + \sum_{s \in \mathcal{S}_l} x_s^{(t)} - c_l \right]_0^{B_l}$$

Comparing with (5), we see that backlogs are related to prices by $b_l(t) = \gamma^{-1} p_l(t)$. Thus, the backlog can be quite large in equilibrium since $\gamma > 0$ is typically small. To remedy this problem, in [9] the price adjustment (5) is modified to

$$p_l^{(t+1)} = [p_l^{(t)} + \alpha b_l^{(t)} + \gamma (\sum_{s \in \mathcal{S}_l} x_s^{(t)} - c_l)]_0^{\infty} \tag{8}$$

where $\alpha > 0$ is a small constant. This extra backlog term will, if equilibrium is achieved, guarantee that $b_l = 0$, i.e. the buffer is cleared. The update algorithm (8) is commonly referred to as REM, whose convergence has not been established yet. In [10], the continuous time version of REM has been shown to be globally stable, but for the original discrete-time case, a proof of stability is available only for the single-link case, see [11]. In this paper, we generalize this proof to the multi-link case using an appropriate Lyapunov function.

To simplify notation, for each source s, we denote the inverse of its marginal utility $U_s'^{-1}$ restricted to $[m_s, M_s]$ as

$$f_s(q_s) := [U_s'^{-1}(q_s)]_{m_s}^{M_s} \tag{9}$$

The overall system is this described by the following set of equations

$$b_l^{(t+1)} = \left[b_l^{(t)} + y_l^{(t)} - c_l \right]_0^{B_l} \tag{10}$$

$$p_l^{(t+1)} = \left[p_l^{(t)} + \alpha b_l^{(t)} + \gamma (y_l^{(t)} - c_l) \right]_0^{\infty} \tag{11}$$

Let (b^*, p^*) be an equilibrium point of this system, and let $q^* = R^T p^*$ be the equilibrium source prices, $x^* = f(q^*)$ the equilibrium source rates, and $y^* = R x^*$ the equilibrium link rates. Clearly, we must have that $b^* = 0$. p^*, on the other hand, need not be zero, indeed its nonzero components correspond to links where $y_l^* = 0$, i.e. where the capacity constraint is active (bottleneck links). It is shown in [6] that the equilibrium is a saddle point of (4); therefore it follows from duality theory that x^* must be the unique global optimum of the primal problem (2)-(3); therefore y^*, q^* are also unique. However, p^* need not be unique, because in general capacity constraints might not be independent. In order to obtain a unique equilibrium price, we make the assumption that the routing matrix R is of full row rank. This assumption guarantees that, for a given vector q of aggregate source prices, there is a unique vector p of source prices satisfying $q = R^T p$.

III. Convergence Analysis

We start our analysis with some simplifying assumptions and definitions. Recall that source functions $f_s(\cdot)$ are strictly decreasing on (m_s, M_s). We make the following assumption on these functions.

Assumption 3.1: *The maximum and minimum rates of decrease for the source function $f_s(\cdot)$ are independent of s, and are given by R, and r, respectively.*

This assumption is not overly restrictive, since for the maximum rate of decrease we can simply pick

$$R = \max_s R_s$$

where $R_s = \max_{p \geq 0} |f_s'(p)|$ is the maximum rate of decrease of source s. Similarly for the minimum rate of decrease we can pick

$$r = \min_s r_s$$

where $r_s = \min_{p \geq 0} |f_s'(p)|$ is the minimum rate of decrease of source s. Note that Assumption 3.1 assumes more than Lipschitz continuity of $f_s(\cdot)$, it also guarantees a minimum rate of decrease, i.e. no flat parts in the graph of $f_s(\cdot)$.

Define Ω as the set of possible states in buffer-length/price space:

$$\Omega = \{(b,p) : b_l \geq 0,\ p_l \geq 0,\ l = 1, \ldots, L\}$$

Note that we omitted the upper bound B_l from the buffer length, because stability of b_l would imply boundedness. This operation, in effect amounts to setting $B_l = \infty$. We assume that the initial condition of the system lies in Ω, i.e. $(b^{(0)}, p^{(0)}) \in \Omega$. This initial condition induces an initial rate vector through (9), where each f_s lies between m_s and M_s. Let us introduce index vectors for links to keep track of the positivity constraints imposed on the states. Define for $(b,p) \in \Omega$,

$$\begin{aligned} I^- &= \{l : p_l + \alpha b_l + \gamma(y_l - c_l) \leq 0\} \\ I^+ &= \{l : p_l + \alpha b_l + \gamma(y_l - c_l) \geq 0\} \\ J^- &= \{l : b_l + y_l - c_l \leq 0\} \\ J^+ &= \{l : b_l + y_l - c_l \geq 0\} \end{aligned}$$

Clearly, these sets are *not* mutually exclusive, a particular link may belong to more than one of these sets. To motivate the analysis for the multi-link case, we first discuss the single-problem in detail.

A. Single-Link Case

Consider the case in which there is a single-link, and a single source[1]. The system then evolves according to

$$\begin{aligned} b^{(t+1)} &= \left[b^{(t)} + x^{(t)} - c\right]_0^\infty \\ p^{(t+1)} &= \left[p^{(t)} + \alpha b^{(t)} + \gamma(x^{(t)} - c)\right]_0^\infty \end{aligned}$$

with the equilibrium point $b^* = 0$, and $x^* = f(p^*) = c$. We first show that $[p^{(t)} + \alpha b^{(t)} + \gamma(x^{(t)} - c)]$ is always nonnegative, thus eliminating the need for the saturation analysis for $p^{(t)}$. Let $(b,p) \in \Omega$, be any point along the trajectories of these updates. We denote the next state by $(\bar{b}, \bar{p})$. Clearly, if $f(p) > c$, i.e. $p < p^*$, we have $p + \alpha b + \gamma(f(p) - c) \geq 0$. Suppose $f(p) \leq c$. Then,

$$p + \alpha b + \gamma(f(p) - c) \geq (p - p^*) + \gamma(f(p) - c)$$

Using Assumption 3.1, if $0 < \gamma < \frac{1}{R}$ we infer that

$$(p - p^*) + \gamma(f(p) - f(p^*)) \geq 0$$

as claimed.

Now, consider a candidate Lyapunov function that is separable in the states (b,p), i.e. $V(b,p) = V_1(b) + V_2(p)$. We will show that such a function cannot be a Lyapunov function for this system. Consider the initial state when p is at equilibrium, $p = p^*$, and $b > 0$. The next state then will be $\bar{b} = b$, and $\bar{p} = p^* + \alpha b$. Thus, the buffer length remains the same, while the link price increases. The first difference of a separable V along the trajectories of the system equals

$$\begin{aligned} \Delta V &= \Delta V_1 + \Delta V_2 = V_1(\bar{b}) - V_1(b) + V_2(\bar{p}) - V_2(p) \\ &= V_1(b) - V_1(b) + V_2(p^* + \alpha b) - V_2(p^*) \\ &= V_2(p^* + \alpha b) > 0 \end{aligned}$$

since $V_2(p)$ must vanish at $p = p^*$, and is positive elsewhere by definition. Hence, we conclude that no separable V would work as a Lyapunov function for the system.

Motivated by the counter-example, we introduce the candidate Lyapunov function:

$$V(b,p) = (p - p^*)^2 + 2M(p - p^*)b + Nb^2$$

where M, N are constants to be determined. This is a quadratic form in $(p - p^*)$ and b, so it is positive definite if $N > 0$, and $N > M^2$. Also, $V(b,p) = 0$ implies $b = 0, p = p^*$, i.e. the function V only vanishes at the equilibrium. Furthermore, being a quadratic V is radially unbounded. We proceed by taking the first difference of $V(b,p)$ along the trajectories of the system:

$$\Delta V(b,p) = V(\bar{b}, \bar{p}) - V(b,p)$$

We need to distinguish between four cases depending on which one of the sets I^-, I^+, J^-, J^+, the single-link l belongs. We have already shown that l cannot belong to the set I^-, which leaves us with two cases.

[1]Multiple sources can be lumped into a single source with an equivalent utility function, see [11]

Case I ($l \in I^+, J^-$): In this case, we have $\bar{p} = p + \alpha b + \gamma(x - c) \geq 0$, and $\bar{b} = 0$, since $b + y - c \leq 0$. Thus,

$$-\frac{1}{\gamma}(p + \alpha b) \leq (x - c) \leq -b \qquad (19)$$

Substituting the state equation into $\Delta V(b,p)$, we obtain

$$\begin{aligned}\Delta V(b,p) &= (p - p^* + \alpha b + \gamma(x-c))^2 - (p-p^*)^2 \\ &\quad -2M(p-p^*)b - Nb^2\end{aligned}$$

Using (19) and Assumption 3.1, we bound $\Delta V(b,p)$ from above as follows:

$$\begin{aligned}\Delta V(b,p) &\leq \alpha^2 b^2 + \gamma^2 R^2 (p-p^*)^2 + 2\alpha b(p-p^*) \\ &\quad -2\alpha\gamma b^2 - 2\gamma r(p-p^*)^2 - 2M(p-p^*)b \\ &\quad -Nb^2 \\ &\leq (\alpha^2 - 2\alpha\gamma - N)b^2 + 2(\alpha - M)(p-p^*)b \\ &\quad +\gamma(\gamma R^2 - 2r)(p-p^*)^2\end{aligned}$$

This is a quadratic in b and $(p - p^*)$, and is nonnegative definite if

$$0 < \gamma < \frac{2r}{R^2}$$

and

$$\alpha^2 - 2\alpha\gamma - N < 0$$

which is satisfied if $0 < \alpha < 2\gamma$. We also require

$$\gamma(\alpha^2 - 2\alpha\gamma - N)(\gamma R^2 - 2r) > 4(\alpha - M)^2$$

to take care of the cross-term $(p - p^*)b$.

Case II ($l \in I^+, J^+$): Now, both $\bar{b}_l$ and $\bar{p}_l$ are nonnegative. Expanding out the $\Delta V(b,p)$ yields

$$\begin{aligned}\Delta V(b,p) &= (p - p^* + \alpha b + \gamma(x-c))^2 - (p-p^*)^2 \\ &\quad +2M(p - p^* + \alpha b + \gamma(x-c))(b + x - c) \\ &\quad -2M(p-p^*)b + N(b + x - c)^2 - Nb^2\end{aligned}$$

We have

$$\begin{aligned}\Delta V(b,p) &\leq \alpha(\alpha + 2M)b^2 \\ &\quad +(\gamma^2 + N + 2\gamma M - 2\gamma r - 2Mr)(x-c)^2 \\ &\quad +2(\alpha M + \gamma M + N + \alpha\gamma + 2\alpha R)(x-c)b\end{aligned}$$

which is negative-definite, if $\alpha < -2M$, $M < 0$, and

$$0 < \gamma < -(M - r) + \sqrt{M^2 - N + r^2}$$

Combining both cases, we conclude that the algorithm is asymptotically stable for the single link case, if the controller gains (α, γ) are chosen appropriately.

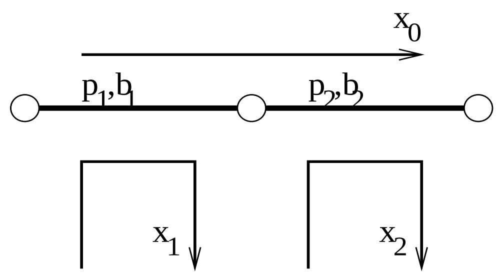

Fig. 1. Two-Link Case

B. Two-Link Case

Consider the two-link, three-source network shown in Figure 1. We will show that a Lyapunov function similar to one given for the single-link case works for two links. For this purpose, let us first consider the direct extension of $V(b,p)$ of the preceding section.

$$V(b,p) = \sum_{l=1}^{2} \left[(p_l - p_l^*)^2 + 2M(p_l - p_l^*)b_l + Nb_l^2\right]$$

As it turns out this function does not monotonically decrease along the trajectories of the system. To obtain a negative drift in $V(b,p)$ we need to consider all cross terms between $(p_1 - p_1^*)$, $(p_2 - p_2^*)$, b_1, and b_2 which results in the following candidate Lyapunov function:

$$\begin{aligned}V(b,p) &= (p_1 - p_1^*)^2 + 2M(p_1 - p_1^*)b_1 + Nb_1^2 \\ &\quad +(p_2 - p_2^*)^2 + 2M(p_2 - p_2^*)b_2 + Nb_2^2 \\ &\quad +2C(p_1 - p_1^*)(p_2 - p_2^*) + 2Db_1 b_2 \\ &\quad +2P(p_1 - p_1^*)b_2 + 2Q(p_2 - p_2^*)b_1\end{aligned}$$

where M, N, P, Q, C, D are constants to be determined. It is not immediately clear that V is positive definite, but conditions can be imposed on the weights M, N, P, Q, C, D so that V is indeed positive definite. Let us first concentrate on the terms

$$V_1 = (p_1 - p_1^*)^2 + (p_2 - p_2^*)^2$$

as it would be representative of the other terms in the sum constituting $V(b,p)$. The first difference of V_1 equals (assuming $1, 2 \in I^+$)

$$\begin{aligned}\Delta V_1 &= (p_1 - p_1^* + \alpha b_1 + \gamma(x_0 + x_1 - c_1))^2 \\ &\quad +(p_2 - p_2^* + \alpha b_2 + \gamma(x_0 + x_2 - c_2))^2 \\ &\quad -(p_1 - p_1^*)^2 - (p_2 - p_2^*)^2\end{aligned}$$

which can be written as

$$\begin{aligned}\Delta V_1 &= (\alpha b_1 + \gamma(x_0 + x_1 - c_1))^2 \\ &\quad +(\alpha b_2 + \gamma(x_0 + x_2 - c_2))^2 \\ &\quad +2\alpha\gamma b_1(x_0 + x_1 - c_1) + 2\alpha\gamma b_2(x_0 + x_2 - c_2) \\ &\quad +2\alpha(p_1 - p_1^*)b_1 + 2\alpha(p_2 - p_2^*)b_2 \\ &\quad +2\gamma(p_1 - p_1^*)(x_0 + x_1 - c_1) \\ &\quad +2\gamma(p_2 - p_2^*)(x_0 + x_2 - c_2)\end{aligned}$$

Now, the last two terms of this sum can be converted into a sum over individual sources, i.e.

$$\begin{aligned} & 2\gamma(p_1 - p_1^*)(x_0 + x_1 - c_1) \\ & +2\gamma(p_2 - p_2^*)(x_0 + x_2 - c_2) \\ = \; & 2\gamma(p_1 - p_1^* + p_2 - p_2^*)(x_0 - x_0^*) \\ & +2\gamma(p_1 - p_1^*)(x_1 - x_1^*) \\ & +2\gamma(p_2 - p_2^*)(x_2 - x_2^*) \\ = \; & 2\gamma(q_0 - q_0^*)(x_0 - x_0^*) \\ & +2\gamma(q_1 - q_1^*)(x_1 - x_1^*) \\ & +2\gamma(q_2 - q_2^*)(x_2 - x_2^*) \end{aligned}$$

where each term is negative since

$$(x_s - x_s^*) = (f_s(q_s) - f_s(q_s^*))$$

and $f_s(\cdot)$ are strictly decreasing. Following along the same lines, other terms of ΔV can be decomposed into sums over sources as well.

As a result, ΔV will consist of two sums of quadratics. The first sum will be over the links (links 1 and 2 in our case), and it will contain quadratic terms of

$$(y_1 - c_1), (y_2 - c_2), (p_1 - p_1^*), (p_2 - p_2^*), b_1, b_2$$

while the second sum will be over the sources (sources $0, 1$, and 2 in our case), and this sum will contain quadratic terms of the variables

$$(x_0 - x_0^*), (x_1 - x_1^*), (x_2 - x_2^*), (q_0 - q_0^*), (q_1 - q_1^*), (q_2 - q_2^*)$$

The overall drift ΔV can in turn be made non-negative definite through an appropriate choice of the controller gains (α, γ). The analysis of this case is rather involved, hence omitted in the conference version of the paper, see [13] for details.

C. Multi-Link Case

For the multi-link case we propose a similar candidate Lyapunov function that contains cross terms of all $(p_l - p_l^*)$'s and b_l's. Nevertheless, the computation of the first difference of this Lyapunov function along the trajectories of the system is rather cumbersome. The stability analysis of the system in the multi-link case is still an open problem.

IV. Modifying REM

In this section, we discuss slight modifications to the REM algorithm that could make stability analysis for multiple-links more tractable. The main difficulty in the Lyapunov-based stability analysis arises from the coupling between the link rates through the price updates. One possibility to avoid this coupling is to add a new state x_s for each source in the network, and update it according to

$$x_s^{(t+1)} = f_s(q_s^{(t)}) \qquad (29)$$

This amounts to updating, both prices (the dual variables), and rates (the primal variables) at the same time to achieve convergence to the global maximum of (2) subject to (3). In [8], this update along with the price update is called the *primal-dual* algorithm, and it can be shown that it is globally asymptotically stable in the continuous-time case. A slightly modified version of (29) is

$$x_s^{(t+1)} = (1 - \beta)x_s^{(t)} + \beta f_s(q_s^{(t)}) \qquad (30)$$

where $0 < \beta < 1$ is a *relaxation* constant.

Either one of the rate updates (29)-(30) can be used in conjunction with the price updates (10), and the buffer-length equation (11) to simplify the stability analysis.

V. Conclusions and Future Work

In this paper, we have discussed stability of REM algorithm for multi-link networks. We have taken a direct Lyapunov approach, and have shown that the gains (α, γ) need to be picked sufficiently small to ensure stability, though picking them too small may result in a slow rate of convergence. An area of future research may be to investigate the optimal choice of these parameters to balance between the conflicting goals of robust stability margins and fast convergence. Several other extensions of this work are possible. For example, the stability of REM in the presence of delay needs to be addressed, as communication/processing/queuing delay is present in any communication network. Also, it may be interesting to investigate stability properties of the algorithm under asynchronous updates.

References

[1] D. P. Bertsekas, R. Gallager Data Networks, Englewood Cliffs, NJ: Prentice-Hall 1987.

[2] W. Stevens TCP/IP Illustrated: The protocols, Volume 1. Addison-Wesley, 1999. 15th printing.

[3] V. Jacobson "Congestion avoidance and control" In *Proceedings of SIGCOMM*, August 1988.

[4] F. P. Kelly "Charging and rate control for elastic traffic" *European Transactions on Telecommunications*, 1997; 8:33-37.

[5] F. P. Kelly, A. Maulloo, D. Tan "Rate control for communication networks: Shadow prices, proportional fairness and stability" *Journal of Operations Research Society*, 1998; 49(3): 237-252.

[6] S. H. Low, D. E. Lapsley "Optimization flow control-I: Basic algorithm and convergence", *IEEE/ACM Transactions on Networking*, 1999; 7(6): 861-874.

[7] S. Kunniyur and R. Srikant. "End-to-end congestion control schemes: utility functions, random losses and ECN marks" In *Proceedings of IEEE Infocom*, March 2000.

[8] Shao Liu, T. Başar, and R. Srikant "Controlling the Internet: A survey and some new results" In *Proceedings of 42nd IEEE Conference on Decision and Control* Mauii, HW, Dec., 2003.

[9] S. Athuraliya, V. H. Li, Steven H. Low and Qinghe Yin "REM: Active Queue Management" *IEEE Transactions on Networking*, Volume:15 Issue:3, pp. 48-53, May/June 2001.

[10] F. Paganini "A Global Stability Result in Network Flow Control", *Systems and Control Letters*; 46(3), pp. 153-163, 2002.

[11] Q. Yin and S. H. Low "Convergence of REM Flow Control at a Single Link", *IEEE Communications Letters*, Volume:5, Issue:3, pp. 119-121, March 2001.

[12] N. V. Yefimov Quadratic Forms and Matrices, Academic Press, NY and London, 1964.

[13] O. C. Imer and T. Başar "Lyapunov functions and stability of network flows," Coordinated Science Laboratory Internal Report 2004.