Chapter 1

DUALITY-BASED TCP CONGESTION CONTROL WITH ERROR ANALYSIS

Mortada Mehyar
California Institute of Technology
morr@caltech.edu

Demetri Spanos
California Institute of Technology
demetri@caltech.edu

Steven H. Low
California Institute of Technology
slow@caltech.edu

Abstract We review optimization models for congestion control, focusing on the dual of a widely studied utility maximization problem. We show that many congestion control algorithms can be modeled within this framework, and conversely, standard techniques for iterative optimization can inspire practical congestion control algorithms. This previous work all assumes precise feedback of congestion information from network to traffic sources. We study the effect of error in feedback information within this duality model. Using standard inexact optimization techniques, we show that, provided relative error is bounded, the network behavior is attracted to a region that contains the optimal point if exact information were available.

1. Introduction

The primary resource of the Internet, bandwidth, is a scarce commodity that must be allocated to individual users according to pre-defined algorithms. The current architecture of the Internet addresses this allocation problem through the interaction of two types of algorithms:

source algorithms, which run on a user's computer, and *link* algorithms, which run on the physical hardware of the network (links and routers).

The function of source algorithms is to update transmission rate in response to perceived congestion of the network. Link algorithms, in contrast, regulate the flow of data by updating (possibly implicitly) some measure of network congestion. This source/link division reflects the inherently distributed design of the Internet, in which no centralized coordination is available.

Transmission Control Protocol (TCP), in its various forms (such as TCP-Reno [6, 14] or TCP-Vegas [3]), is the most common source algorithm on the current Internet. Similarly, many link algorithms exist, such as DropTail, Random Early Detection (RED) [5], and Random Exponential Marking (REM) [1]. Collectively, all these algorithms fall under the general headings of Congestion Control (for source algorithms) and Active Queue Management (for link algorithms).

The interaction of these algorithms is generally very complex, and many models have been developed for their study. In particular it has been shown that several important features of this interaction can be understood as the execution of a distributed network optimization procedure. The underlying optimization problem, first presented as a model for Internet congestion control in [7], seeks an optimal allocation of transmission rates subject to bandwidth (capacity) constraints on each link.

In this article, we first review results pertaining to the dual of the aforementioned network optimization problem [8–10]. We then address an issue that arises in the application of these theoretical results in real networks (see [12] for preliminary results). Whereas the primal problem is concerned with optimal selection of transmission rates (primal variables), the dual problem addresses selection of the congestion measures on each link (this is also called the "price").

Section 1.2 shows that a very general class of source and link algorithms can be viewed as a (primal-dual) algorithm for simultaneous solution of the transmission rate and link price optimization. The main result of this section is that any equilibrium of the network (which is determined by the specific choice of source and link algorithms) is a primal-dual optimum. This allows one to understand equilibrium properties of quite general networks in terms of optima of a single optimization problem.

The work described in Section 1.3 reverses the approach of the previous section: rather than taking the source and link algorithms as given, it begins with the optimization problem and its dual. We show that the dual has a very interesting distributed structure, and that standard iterative techniques for solving optimization problems naturally inspire source and link algorithms (in particular, dual-gradient and Newton-like algorithms). We thus see that this optimization framework for network modeling can be used either as a tool for understanding given algorithms, *or* as a guide for new source and link algorithms.

All the models reviewed above assume that precise link prices are fed back to sources for their rate adjustment. Section 1.4 addresses the fact that sources on real networks do not have direct access to congestion prices, and can only estimate them based on locally observable quantities, such as packet losses or queuing delay.

We show that this non-ideal behavior at the sources can be viewed as an inexact gradient calculation. We then apply standard techniques for analysis of inexact gradient methods to characterize the behavior of the theoretically-inspired algorithms in the presence of this more realistic price-estimation scheme. We prove that as long as the relative error is bounded, the optimization flow control scheme will still "converge" in the sense that it will drive the link utilization to an attraction region. The core of the argument is that reduction of the dual function can still be achieved in the presence of inexact gradient calculation. The analysis also suggests an optimal choice of stepsize that guarantees the greatest decrease in the dual function at each iteration. Finally, we apply these results to typical types of error in price feedback and illustrate the attraction region and stepsize optimality with numerical examples.

2. Duality Model

A network is modeled as a set L of links (scarce resources) with finite capacities $c = (c_l, l \in L)$. They are shared by a set S of sources indexed by s. Each source s uses a set $L_s \subseteq L$ of links. The sets L_s define an $L \times S$ routing matrix[1]

$$R_{ls} = \begin{cases} 1 & \text{if } l \in L_s \\ 0 & \text{otherwise} \end{cases}$$

Associated with each source s is its transmission rate $x_s(t)$ at time t, in packets/sec. Associated with each link l is a scalar congestion measure $p_l(t) \geq 0$ at time t. Following the notation of [13], let $y_l(t) = \sum_s R_{ls} x_s(t)$ be the aggregate source rate at link l and let $q_s(t) = \sum_l R_{ls} p_l(t)$ be the end-to-end congestion measure for source s. In vector notation, we have

[1]We abuse notation to use L and S to denote sets and their cardinality.

($\cdot^T$ denotes transpose)

$$\begin{aligned} y(t) &= Rx(t) \\ q(t) &= R^T p(t) \end{aligned}$$

Here, $x(t) = (x_s(t), s \in S)$ and $q(t) = (q_s(t), s \in S)$ are in $\Re_+^{|S|}$, and $y(t) = (y_l(t), l \in L)$ and $p(t) = (p_l(t), l \in L)$ are in $\Re_+^{|L|}$ ($\Re_+$ denotes non-negative real). Source s can observe its own rate $x_s(t)$ and the end-to-end congestion measure $q_s(t)$ of its path, but not the vector $x(t)$ or $p(t)$, nor other components of $q(t)$. Similarly, link l can observe just local congestion $p_l(t)$ and flow rate $y_l(t)$.

The source rate $x_s(t)$ is adjusted in each period according to a function F_s based only on $x_s(t)$ and $q_s(t)$: for all s,

$$x_s(t+1) = F_s(x_s(t), q_s(t)) \tag{1.1}$$

The link congestion measure $p_l(t)$ is adjusted in each period based only on $p_l(t)$ and $y_l(t)$, and possibly some internal (vector) variable $v_l(t)$, such as the queue length at link l. This can be modeled by some functions (G_l, H_l): for all l,

$$p_l(t+1) = G_l(y_l(t), p_l(t), v_l(t)) \tag{1.2}$$
$$v_l(t+1) = H_l(y_l(t), p_l(t), v_l(t)) \tag{1.3}$$

where G_l is non-negative so that $p_l(t) \geq 0$. Here, F_s models TCP algorithms (e.g., Reno or Vegas) and (G_l, H_l) model AQM's (e.g., RED, REM); see the next section. We will often refer to AQM's by G_l, without explicit reference to the internal variable $v_l(t)$ or its adaptation H_l.

We assume that (1.1)–(1.3) has a set of equilibria (x, p). The fixed point of (1.1) defines an implicit relation between equilibrium rate x_s and end-to-end congestion measure q_s:

$$x_s = F_s(x_s, q_s)$$

Assume F_s is continuously differentiable and $\partial F_s / \partial q_s \neq 0$ in the open set $A := \{(x_s, q_s) | x_s > 0, q_s > 0\}$. Then, by the implicit function theorem, there exists a unique continuously differentiable function f_s from $\{x_s > 0\}$ to $\{q_s > 0\}$ such that

$$q_s = f_s(x_s) > 0 \tag{1.4}$$

To extend the mapping between x_s and q_s to the closure of A, define

$$f_s(0) = \inf \{q_s \geq 0 \mid F_s(0, q_s) = 0\} \tag{1.5}$$

possibly ∞. If $(x_s, 0)$ is an equilibrium point, $F_s(x_s, 0) = x_s$, then define

$$f_s(x_s) \quad = \quad 0 \tag{1.6}$$

Define the utility function of each source s as

$$U_s(x_s) \quad = \quad \int f_s(x_s) dx_s, \qquad x_s \geq 0 \tag{1.7}$$

that is unique up to a constant.

Being an integral, U_s is a continuous function. Since $f_s(x_s) = q_s \geq 0$ for all x_s, U_s is nondecreasing. It is reasonable to assume that f_s is a nonincreasing function – the more severe the congestion, the smaller the rate. This implies that U_s is concave. If f_s is *strictly* decreasing, then U_s is strictly concave since $U_s''(x_s) < 0$. An increasing utility function implies a greedy source – a larger rate yields a higher utility – and concavity implies diminishing return.

Now consider the problem of maximizing aggregate utility formulated in [7]:

$$\max_{x \geq 0} \qquad \sum_s U_s(x_s) \qquad \text{subject to } Rx \leq c \tag{1.8}$$

The constraint says that, at each link l, the flow rate y_l does not exceed the capacity c_l. An optimal rate vector x^* exists since the objective function in (1.8) is continuous and the feasible solution set is compact. It is unique if U_s are *strictly* concave. As the sources are coupled through the shared links (the capacity constraint), solving for x^* directly, however, may require coordination among possibly all sources, and hence is infeasible in a large network. The key to understanding the equilibrium of (1.1)–(1.3) is to *regard $x(t)$ as primal variables, $p(t)$ as dual variables, and $(F, G) = (F_s, G_l, s \in S, l \in L)$ as a distributed primal-dual algorithm to solve the primal problem (1.8) and its Lagrangian dual* (see [8]):

$$\min_{p \geq 0} \qquad \sum_s \max_{x_s \geq 0} (U_s(x_s) - x_s q_s) + \sum_l p_l c_l \tag{1.9}$$

Hence, the dual variable is a precise measure of congestion in the network. The dual problem has an optimal solution since the primal problem is feasible.

We will interpret the equilibria (x^*, p^*) of (1.1)–(1.3) as solutions of the primal and dual problem, and that (F, G) iterates on both the primal and dual variables together in an attempt to solve both problems.

We summarize the assumptions on (F, G, H):

C1: For all $s \in S$ and $l \in L$, F_s and G_l are non-negative functions. Moreover, equilibrium points of (1.1)–(1.3) exist.

C2: For all $s \in S$, F_s are continuously differentiable and $\partial F_s/\partial q_s \neq 0$ in $\{(x_s, q_s)|x_s > 0, q_s > 0\}$; moreover, f_s in (1.4) are nonincreasing.

C3: If $p_l = G_l(y_l, p_l, v_l)$ and $v_l = H_l(y_l, p_l, v_l)$, then $y_l \leq c_l$ with equality if $p_l > 0$.

C4: For all $s \in S$, f_s are strictly decreasing.

Condition C1 guarantees that $(x(t), p(t)) \geq 0$ and $(x^*, p^*) \geq 0$. C2 guarantees the existence and concavity of utility function U_s. C3 guarantees the primal feasibility and complementary slackness of (x^*, p^*). Finally condition C4 guarantees the uniqueness of optimal x^*. The following theorem is proved in [9].

THEOREM 1.1 *Suppose assumptions C1 and C2 hold. Let (x^*, p^*) be an equilibrium of (1.1)–(1.3). Then (x^*, p^*) solves the primal problem (1.8) and the dual problem (1.9) with utility function given by (1.7) if and only if C3 holds. Moreover, if assumption C4 holds as well, then U_s are strictly concave and the optimal rate vector x^* is unique.*

Proof. The discussion after the definition (1.7) of U_s proves the second claim when C4 holds, so we only prove the first claim.

By duality theory (e.g., [2]), (x^*, p^*) is primal-dual optimal if and only if x^* is primal feasible, p^* is dual feasible, complementary slackness holds, and

$$x^* = \arg\max_{x \geq 0} L(x, p^*) \tag{1.10}$$

where L is the Lagrangian of (1.8) defined as:

$$L(x, p) = \sum_s U_s(x_s) + \sum_l p_l(c_l - \sum_s R_{ls} x_s)$$

Hence, to prove the first claim, we only need to establish (1.10). Now

$$\begin{aligned}
&\max_{x \geq 0} L(x, p^*) \\
&= \max_{x \geq 0} \sum_s U_s(x_s) + \sum_l p_l^* \left(c_l - \sum_s R_{ls} x_s \right) \\
&= \sum_s \max_{x_s \geq 0} \left(U_s(x_s) - x_s \sum_l R_{ls} p_l^* \right) + \sum_l p_l^* c_l
\end{aligned}$$

By construction of U_s, we have from (1.7) and (1.4) that, for any equilibrium at which $x_s^* > 0$, (x^*, p^*),

$$U_s'(x_s^*) = f_s(x_s^*) \quad = \quad q_s^* = \sum_l R_{ls} p_l^* \tag{1.11}$$

Note that if $q_s^* = 0$, then (1.11) holds by (1.6). If $x_s^* = 0$, we have from (1.5)

$$U_s'(0) \ = f_s(0) \quad \leq \quad q_s^* \tag{1.12}$$

But, (1.11)–(1.12) implies that

$$\frac{\partial L}{\partial x_s}(x^*, p^*) \quad \leq \quad 0$$

with equality if $x_s^* > 0$. Since $L(x, p^*)$ is concave in x, this is the necessary and sufficient Karush-Kuhn-Tucker condition for x^* to maximize $L(x, p^*)$ over $x \geq 0$. Hence the proof is complete. ■

In Table 1.1, we give some examples of specific source and link algorithms. The parameter D_s represents the round-trip-time of a packet. In Vegas, α_s is a scalar parameter and d_s is the propagation delay. Finally, the various quantities appearing in the AQM table are internal variables and system parameters, which we do not discuss here; see [9] for details.

It is prudent to summarize what we have done thus far. We have examined a general class of update algorithms for the rates, prices, and internal variables. We then showed how to construct an optimization problem of the form presented in [7], and that equilibria of the source and link algorithms correspond to optima.

Note that we had to assume existence of equilibria in order to obtain this formalism. Certainly, there is no guarantee that the algorithms we have modeled have equilibria. On the other hand, the assumptions we have made about the algorithms are so general that one cannot really expect a guarantee of an equilibrium.

Nonetheless, we would like to understand at least *some* algorithms which have a provable (and stable) equilibrium at the optimum. This is the motivation of the next section, in which we begin with the formal optimization problem, and demonstrate source and link algorithms (based on a gradient-projection scheme) which provably converge to the primal-dual optimal point. Perhaps more importantly, the gradient-projection scheme gives rise to *distributed* source and link algorithms, and in this regard can serve as an inspiration for practical algorithms on a real network.

TCP		
Reno-1	$F_s(x_s(t), q_s(t))$	$\left[x_s(t) + \frac{1-q_s(t)}{D_s^2} - \frac{2}{3}q_s(t)x_s^2(t)\right]^+$
	Utility	$\frac{\sqrt{3/2}}{D_s} \tan^{-1}\left(\sqrt{\frac{2}{3}}\, x_s D_s\right)$
Reno-2	$F_s(x_s(t), q_s(t))$	$\left[x_s(t) + \frac{1-x_s(t)D_s q_s(t)}{D_s^2} - \frac{2}{3}q_s(t)x_s^2(t)\right]^+$
	Utility	$\frac{1}{D_s} \log \frac{x_s D_s}{2x_s D_s + 3}$
Vegas	$F_s(x_s(t), q_s(t))$	$\begin{cases} x_s(t) + \frac{1}{D_s^2} & \text{if } x_s(t) < \overline{x}_s(t) \\ x_s(t) - \frac{1}{D_s^2} & \text{if } x_s(t) > \overline{x}_s(t) \\ x_s(t) & \text{otherwise} \end{cases}$
	Utility	$\alpha_s d_s \log x_s$
AQM		
RED	$G_l(y_l(t), p_l(t), v_l(t))$	$\begin{cases} 0 & r_l(t+1) \leq \underline{b}_l \\ \rho_1(r_l(t+1) - \underline{b}_l) & \underline{b}_l \leq r_l(t+1) \leq \overline{b}_l \\ \rho_2(r_l(t+1) - \overline{b}_l) + m_l & \overline{b}_l \leq r_l(t+1) \leq 2\overline{b}_l \\ 1 & r_l(t+1) \geq 2\overline{b}_l \end{cases}$
	$H_l(y_l(t), p_l(t), v_l(t))$	$b_l(t+1) = [b_l(t) + y_l(t) - c_l]^+$ $r_l(t+1) = (1-\alpha_l)r_l(t) + \alpha_l b_l(t)$
REM	$G_l(y_l(t), p_l(t), v_l(t))$	$1 - \phi^{r_l(t+1)}$
	$H_l(y_l(t), p_l(t), v_l(t))$	$b_l(t+1) = [b_l(t) + y_l(t) - c_l]^+$ $r_l(t+1) = [r_l(t) + \gamma(\alpha_l b_l(t) + y_l(t) - c_l)]^+$
Delay	$G_l(y_l(t), p_l(t), v_l(t))$	$p_l(t+1) = [p_l(t) + \frac{y_l(t)}{c_l} - 1]^+$

Table 1.1. A list of some common TCP/AQM algorithms and their associated functions. The notation $[z]^+$ signifies $\max\{z, 0\}$.

3. Optimization Flow Control

We will now discuss the construction of flow control algorithms based directly on the network optimization problem. That is, we view the optimization formalism as representing some kind of design goal, and seek source and link algorithms which provably converge to the optimum.

We will demonstrate a dual-gradient algorithm for solving this optimization, which results in corresponding source and link algorithms that are inherently distributed. We will show global convergence of this algorithm in a synchronous implementation. Analogous asynchronous results can also be obtained, but we will omit them as they are quite technical.

We now introduce some additional notation and assumptions. Recall that for each source s, q_s (the sth component of $p^T R$) is the path bandwidth price that s faces. Let $x_s(p)$ be the unique maximizer in (1.10). We will abuse notation and use $x_s(\cdot)$ both as a function of scalar price $q_s \in \Re_+$ and of vector price $p \in \Re_+^{|L|}$. When the argument of $x_s(\cdot)$ is a scalar, by the Karush–Kuhn–Tucker theorem, $x_s(q_s)$ is given by

$$x_s(p) \quad = \quad [U_s'^{-1}(q_s)]_{m_s}^{M_s} \tag{1.13}$$

where $[z]_a^b = \min\{\max\{z, a\}, b\}$. Here $U_s'^{-1}$ is the inverse of U_s', which exists over the range $[U_s'(M_s), U_s'(m_s)]$ since U_s' is continuous and U_s *strictly* concave (condition C1 below).

When the argument of $x_s(\cdot)$ is a vector, $x_s(p) = x_s(q_s)$. The meaning should be clear from the context. Also, $x(p) = (x_s(q_s), s \in S)$.

We make the following assumptions regarding the utility functions:

A1: On the interval $I_s = [m_s, M_s]$, the utility functions U_s are increasing, strictly concave, and twice continuously differentiable. For feasibility, assume $\sum_{s \in S(l)} m_s \le c_l$ for all l.

A2: The curvatures of U_s are bounded away from zero on I_s: $-U_s''(x_s) \ge 1/\overline{\alpha}_s > 0$ for all $x_s \in I_s$.

We will make use of the notations $\overline{L} := \max_{s \in S} |L(s)|$, $\overline{S} := \max_{l \in L} |S(l)|$, and $\overline{\alpha} := \max\ \{\overline{\alpha}_s,\ s \in S\}$.

3.1 Synchronous Distributed Algorithm

In this section we present the basic source and link algorithms and prove its convergence under conditions A1 and A2. From now on we assume that the algorithms are synchronous (i.e., all sources update at the same time, and so do the links). The asynchronous version of this algorithm can be shown to converge to the optimum as well and the proof can be found in [8].

We will solve the dual problem using gradient projection method (e.g., [11, 4]) where link prices are adjusted in opposite direction to the gradient $\nabla D(p)$:

$$p_l(t+1) \quad = \quad [p_l(t) - \gamma \frac{\partial D}{\partial p_l}(p(t))]^+ \tag{1.14}$$

Here $\gamma > 0$ is a stepsize, and $[z]^+ = \max\{z, 0\}$. Recall that $x_s(p)$ is given by (1.13). Then

$$D(p) \quad = \quad \sum_s (U_s(x_s(p)) - x_s(p)q_s) + \sum_l p_l c_l.$$

Since U_s are *strictly* concave, $D(p)$ is continuously differentiable ([4]) with derivatives given by

$$\frac{\partial D}{\partial p_l}(p) \quad = \quad c_l - y_l \tag{1.15}$$

Substituting (1.15) into (1.14) we obtain the following price adjustment rule for link $l \in L$:

$$p_l(t+1) \quad = \quad [p_l(t) + \gamma(y_l(t) - c_l)]^+ \tag{1.16}$$

The decentralized nature of is striking: though the dual problem is not separable in p, given aggregate source rate y_l that goes through link l, the adjustment algorithm (1.16) is completely distributed and can be implemented by individual links using only local information. We summarize:

Algorithm: Synchronous Gradient Projection
Link l's algorithm:
At times $t = 1, 2, \ldots,$ link l:

1 *Receives rates $x_s(t)$ from all sources $s \in S(l)$ that go through link l.*

2 *Computes a new price*

$$p_l(t+1) \quad = \quad [p_l(t) + \gamma(y_l(t) - c_l)]^+$$

3 *Communicates new price $p_l(t+1)$ to all sources $s \in S(l)$ that use link l.*

Source s's algorithm:
At times $t = 1, 2, \ldots,$ source s:

1. *Receives from the network the sum of link prices in its path $q_s(t)$.*
2. *Chooses a new transmission rate $x_s(t+1)$ for the next period:*

$$x_s(t+1) \quad = \quad x_s(q_s(t))$$

3. *Communicates new rate $x_s(t+1)$ to links $l \in L(s)$ in its path.*

The convergence of this algorithm is shown in [8].

THEOREM 1.2 *Suppose assumptions A1–A2 hold and the stepsize satisfies $0 < \gamma < 2/\overline{\alpha}\overline{LS}$. Then starting from any initial rates $m \leq x(0) \leq M$ and prices $p(0) \geq 0$, every accumulation point (x^*, p^*) of the sequence $(x(t), p(t))$ generated by Algorithm A1 is primal–dual optimal.*

Proof (Sketch). Assumptions A1 and A2 can be shown to imply that the following Lipschitz condition holds

$$||\nabla D(p) - \nabla D(p')||_2 \quad \leq \quad \overline{\alpha}\overline{LS}\ ||p - p'||_2$$

for all $p, p' \geq 0$ [8]. This is a sufficient condition for convergence of the dual-gradient algorithm [2], and convergence to a primal–dual optimal point follows from the concavity assumptions. ■

We have thus accomplished our goal of constructing source and link algorithms which provably converge to the optimum. The fact that these algorithms are also naturally distributed makes the result even more interesting, as it suggests applications in realistic algorithms on real networks.

4. Optimization Flow Control with Estimation Error

One practical drawback of the proposed dual-gradient method is the reliance on explicit communication of price information. Schemes such as RED [5] and REM [1] use a congestion-based queue-management protocol which, in the above context, amounts to an implicit price-notification scheme. Although this mechanism is more practical than the explicit transmission of price information, it suffers from various errors inherent in the implicit price-notification.

One particular source of error inherent to any physical implementation is the limited information available to individual sources, i.e., the receipt of acknowledgments and the round-trip-time (RTT) for each packet transmitted. The prices (congestion measures) in these two cases are, respectively, loss probability and queueing delay. The exact price is either very hard to estimate (loss probability) or very noisy (queueing delay). Our aim in this section is to understand the effects of such errors on the performance of the dual-gradient algorithm [12].

4.1 Price Estimation Error as Inexact Gradient

It is evident that the price update in the algorithm of Section 1.3 is dependent on the rate update. Thus, an erroneous rate update will result in a corresponding error in the price update. In particular, the direction of the price update will, in general, not be along the gradient direction, but along some perturbed direction. *Thus, the effect of inexact price estimation at the sources amounts to an inexact calculation of the gradient at the links.*

The advantage of the inexact gradient viewpoint is that it allows us to embed the phenomenon of price estimation error in the optimization flow control framework. We will show that in the presence of error, the above algorithm will still drive the system to a region around the optimum, under a slight modification of the stepsize bound.

4.2 Attraction under Inexact Gradient

In this section we will characterize the steady-state dynamics of link utilization in terms of an *attraction region*, which we define as follows:

Definition: A set $A_l \subset \Re_+$ is called an attraction region for link l if there exists an integer N such that for all initial conditions (source rates and link prices), $y_l(n) \in A_l$ for some n less than N.

We remark on two important subtleties. First, this definition does not require that the trajectory remain within the attraction region after entering. It is thus not required to be an invariant set. Second, this *does* imply that if the trajectory ever leaves the attraction region, it will return to the attraction region within N steps.

We will show that as long as the relative error is bounded, the optimization flow control scheme will still "converge" in the sense that it will drive the link utilization to an attraction region. The core of the following argument is that reduction of the dual function can still be achieved in the presence of inexact gradient calculation.

At each time t, the lth component of the exact gradient is given by

$$g_l(t) = c_l - \sum_{s \in S(l)} U_s^{'-1}(q_s(t))$$

Let $v(t)_s$ be the estimation error at each source, and define the estimated price at each source by $\tilde{q}_s(t) = q_s(t) + v_s(t)$. Hence, the rate update is $x_s(t) = U_s'^{-1}(\tilde{q}_s(t))$. Thus, the inexact gradient link l actually uses is

$$\tilde{g}_l(t) = c_l - \sum_{s \in S(l)} U_s^{'-1}(\tilde{q}_s(t))$$

The error in the lth component of the gradient is therefore bounded by

$$\begin{aligned} |\tilde{g}_l(t) - g_l(t)| &= \left| \sum_{s \in S(l)} [U_s^{'-1}(q_s(t)) - U_s^{'-1}(\tilde{q}_s(t))] \right| \\ &\leq \sum_{s \in S(l)} \left| [U_s^{'-1}(q_s(t)) - U_s^{'-1}(\tilde{q}_s(t))] \right| \\ &\leq \sum_{s \in S(l)} \alpha_s |v_s(t)| \end{aligned}$$

where $1/\alpha_s$ is the lower bound on the curvature of $U_s(x)$ and therefore α_s is a global Lipschitz constant for $U_s^{'-1}(q_s(t))$ by the Mean Value Theorem.

The following is a sufficient condition which guarantees that the inexact gradient will still be in a descent direction:

$$\sum_{s \in S(l)} \alpha_s |v_s(t)| \leq \eta |c_l - y_l(t)|, \forall l. \tag{1.17}$$

where $0 \leq \eta < 1$ can be thought of as the relative error. This condition simply ensures that the error is not large enough to completely negate the gradient, and so the dual function can still be reduced in the direction of the inexact gradient.

When inequality (1.17) is not satisfied, no conclusion can be drawn, as the above condition is merely sufficient for convergence. Nonetheless, we can show that the region where (1.17) fails, i.e., where the following holds

$$\sum_{s \in S(l)} \alpha_s |v_s(t)| > \eta |c_l - y_l(t)|, \text{for some } l \tag{1.18}$$

contains an attraction region.

THEOREM 1.3 *The solution set of (1.18) is an attraction region, provided*

$$0 < \gamma < \frac{2}{\bar{\alpha}\bar{L}\bar{S}}(1-\eta) \tag{1.19}$$

Proof. The condition in the definition of an attraction region will be verified in two steps:

a) Choice of Stepsize Since $\nabla D = g$ (the exact gradient) is Lipschitz with a Lipschitz constant $\bar{\alpha}\bar{L}\bar{S}$ [8], the Descent Lemma ([2] proposition A.24) implies

$$D(p-\gamma\tilde{g}) \leq D(p) - \gamma\langle g, \tilde{g}\rangle + \frac{\bar{\alpha}\bar{L}\bar{S}}{2}\gamma^2\|\tilde{g}\|^2 \tag{1.20}$$

where $\langle g, \tilde{g}\rangle$ is the Euclidean inner product and $\|g\|$ is the Euclidean norm. Then we see

$$0 < \gamma < \frac{2}{\bar{\alpha}\bar{L}\bar{S}}\frac{\langle g, \tilde{g}\rangle}{\|\tilde{g}\|^2}$$

guarantees that the change in the dual function

$$\Delta D := D(p-\gamma\tilde{g}) - D(p)$$

is strictly less than 0. Therefore when (1.17) holds, the best bound on γ that guarantees descent is the solution of the following optimization problem:

$$\begin{array}{cc} \min_g & \frac{2}{\bar{\alpha}\bar{L}\bar{S}}\frac{\langle g, \tilde{g}\rangle}{\|\tilde{g}\|^2} \\ \text{subject to} & \|\tilde{g}_l - g_l\| \leq \eta\|\tilde{g}_l\|, \forall l \end{array}$$

Since $\langle g, \tilde{g}\rangle = \sum_l g_l\tilde{g}_l$, it is easy to see that the minimum occurs at point O where

$$\langle g, \tilde{g}\rangle \quad = \quad (1-\eta)\|\tilde{g}\|^2 \tag{1.21}$$

and therefore the minimum bound for γ that guarantees descent is $\frac{2}{\bar{\alpha}\bar{L}\bar{S}}(1-\eta)$.

b) Entry in Finite Steps From (1.20) we see that the minimal decrease of the dual function in each step is

$$\begin{aligned}\Delta D &\leq \frac{1}{2}\bar{\alpha}\bar{L}\bar{S}\|\tilde{g}\|^2\gamma\left(\gamma - \frac{\langle g,\tilde{g}\rangle}{\|\tilde{g}\|^2}\frac{2}{\bar{\alpha}\bar{L}\bar{S}}\right)\\ &\leq \frac{1}{2}\bar{\alpha}\bar{L}\bar{S}\|\tilde{g}\|^2\gamma\left[\gamma - (1-\eta)\frac{2}{\bar{\alpha}\bar{L}\bar{S}}\right]\\ &\leq \frac{1}{2}\bar{\alpha}\bar{L}\bar{S}\|\tilde{g}\|^2\left[\gamma^2 - (1-\eta)\frac{2\gamma}{\bar{\alpha}\bar{L}\bar{S}}\right] \qquad (1.22)\end{aligned}$$

where $\|\tilde{g}\|^2$ is *strictly* positive since (1.17) holds and $|v_s(t)|$ is in general not identically zero (else there would be no estimation error). Therefore as long as $0 < \gamma < \frac{2}{\bar{\alpha}\bar{L}\bar{S}}(1-\eta)$, the dual function is decreased by a *finite* amount in each iteration. Now since the primal problem is, by hypothesis, feasible, the dual function is lower bounded [2]. Therefore the inequality (1.17) must fail after a *finite* number of steps, or it would contradict the fact that the dual function is lower bounded. In other words, (1.18) must hold after a finite number of steps, i.e., the trajectory of $y_l(t)$ enters the solution set of (1.18) after a finite number of steps. ■

Some comments on the relationship between η and γ are now in order. First, note that larger values of η will result in smaller solution sets for (1.18). Of course, in order to guarantee that this is an attraction region, (1.19) must be satisfied. So, a larger η corresponds to a tighter attraction region, but demands a smaller γ. Conversely, given a choice of γ, (1.19) constrains the maximal η consistent with the above analysis, and hence the smallest obtainable attraction region.

The above proof demonstrates that convergence can be guaranteed with *any* γ satisfying (1.19), but like the convergence proof in [8], it does not suggest a criterion for selecting γ within this range. As is usual in iterative optimization algorithms, there is a tradeoff between taking larger steps at each iteration (i.e., selecting a large γ) and ensuring that the (inexact) gradient remains a good predictor of local function behavior (i.e., selecting a small γ).

It turns out that we can obtain a satisfactory optimality result which strongly suggests choosing

$$\gamma_{opt} := \frac{(1-\eta)}{\bar{\alpha}\bar{L}\bar{S}}$$

which happens to be half of the bound imposed by the convergence criterion.

THEOREM 1.4 *The choice of stepsize γ_{opt} has the following properties:*

a) Worst-case optimality *It is the worst-case optimal γ, in the sense that it maximizes the Lipschitz-bounded progress in the dual function at each iteration.*

b) Superiority to smaller γ *At each iteration, γ_{opt} generates a new price with a smaller value of the dual function than any smaller choice of γ.*

Proof.
a) From (1.22), the worst-case progress of the dual function is

$$\frac{\bar{\alpha}\bar{L}\bar{S}}{2}\|\tilde{g}\|^2\left[\gamma^2 - \frac{2\gamma}{\bar{\alpha}\bar{L}\bar{S}}(1-\eta)\right]$$

A simple calculation shows that the minimum of the quadratic as a function of γ occurs when $\gamma = \gamma_{opt}$.

b) Consider the directional derivative $\frac{d}{d\gamma}D(p - \gamma\tilde{g}) = -\langle\nabla D(p - \gamma\tilde{g}), \tilde{g}\rangle$. The Lipschitz property of ∇D implies:

$$\|\nabla D(p - \gamma\tilde{g}) - \nabla D(p)\| \leq \bar{\alpha}\bar{L}\bar{S}\gamma\|\tilde{g}\|$$

Therefore the magnitude of the difference between the directional derivatives at p and $p - \gamma\tilde{g}$ is

$$\|\langle\nabla D(p - \gamma\tilde{g}) - \nabla D(p), \tilde{g}\rangle\| \leq \bar{\alpha}\bar{L}\bar{S}\gamma\|\tilde{g}\|^2$$

Here we have used the Cauchy-Schwarz inequality and the Lipschitz bound. This in turn implies:

$$\begin{aligned}\frac{d}{d\gamma}D(p - \gamma\tilde{g}) &\leq -\langle\nabla D(p), \tilde{g}\rangle + \bar{\alpha}\bar{L}\bar{S}\gamma\|\tilde{g}\|^2 \\ &= -\langle g, \tilde{g}\rangle + \bar{\alpha}\bar{L}\bar{S}\gamma\|\tilde{g}\|^2 \\ &\leq \left[-(1-\eta) + \bar{\alpha}\bar{L}\bar{S}\gamma\right]\|\tilde{g}\|^2 \\ &= (\gamma - \gamma_{opt})\bar{\alpha}\bar{L}\bar{S}\|\tilde{g}\|^2\end{aligned}$$

In the second last line we have applied (1.21), since we are only interested in points where the algorithm has not driven the system to the attraction region and can hence provably decrease the dual function at each iteration.

Thus, whenever γ is chosen to be smaller than γ_{opt}, the derivative of the dual function with respect to γ is strictly negative. This implies that

γ_{opt} achieves greater decrease in the dual function at each iteration than any smaller γ.

■

4.3 Quantization Error of Marking

Theorems 1.3 and 1.4 give us some general understanding of link utilization and stepsize selection in the presence of price estimation error. Note that we have not made any assumption on what the error process $v_s(t)$ is. We will now try to model $v_s(t)$ by looking at some typical situations where errors occur.

When the congestion measure is loss probability, e.g., when routers implement RED or REM, during each RTT the source s sends out one window size w_s of packets and has to estimate q_s by observing how many packets are dropped or marked. The fraction of packets lost is an instantaneous estimator of q_s and is subject to two kinds of errors: quantization and probabilistic fluctuation. For example if $w_s = 4$, then $\tilde{q_s} \in \{0, 0.25, 0.5, 0.75, 1\}$. So, if the actual price occurs at some intermediate value, say $q_s = .3$, the closest one could estimate would be $\tilde{q_s} = .25$. We call this the quantization error. Further, due to the probabilistic nature of the dropping scheme, we could get (albeit with lower probability), say $\tilde{q_s} = 1$ as the estimate of q_s and incur a larger error in the effective gradient. We call this the fluctuation error.

It can be seen that, if only the quantization error is present, $|v_s(t)|$ will be bounded by $\frac{1}{2w_s(t)} = \frac{1}{2d_s x_s(t)}$, where d_s is the RTT of source s which is assumed to be constant. Using this specific error model, condition (1.18) becomes

$$\sum_{s\in S(l)} \frac{\alpha_s}{2\eta d_s x_s(t)} > \left| c_l - \sum_{s\in S(l)} x_s(t) \right|.$$

In the single-source-single-link case this reduces to

$$\frac{\alpha}{2\eta d x(t)} > |c - x(t)|$$

and we can solve the inequality and see the attraction region is given by

$$\left(\frac{c}{2} + \frac{1}{2}\sqrt{c^2 - \frac{2\alpha}{\eta d}}, \frac{c}{2} + \frac{1}{2}\sqrt{c^2 + \frac{2\alpha}{\eta d}} \right)$$

Similarly, we can find the following attraction region for a general network (as in [12]):

$$\left(\frac{c_l}{2}+\frac{1}{2}\sqrt{c_l^2-b_l|S(l)|^2\frac{(M^l+m^l)^2}{4M^lm^l}},\frac{c_l}{2}+\frac{1}{2}\sqrt{c_l^2+b_l|S(l)|^2\frac{(M^l)+m^l)^2}{4M^lm^l}}\right)$$

where $m^l = \min_{s(l)} m_s$, $M^l = \max_{s(l)} M_s$, and $|S(l)|$ is the number of sources sharing link l.

To illustrate the effects of quantization error, we present some simple simulation results. We use quadratic utility functions (which conveniently result in linear network dynamics). Our first example is a two-source one-link network.

Figure 1.1 shows the resulting dynamics of the two source rates, as well as the aggregate rate on the link. We observe that the aggregate rate dynamics is nicely characterized by the attraction region.

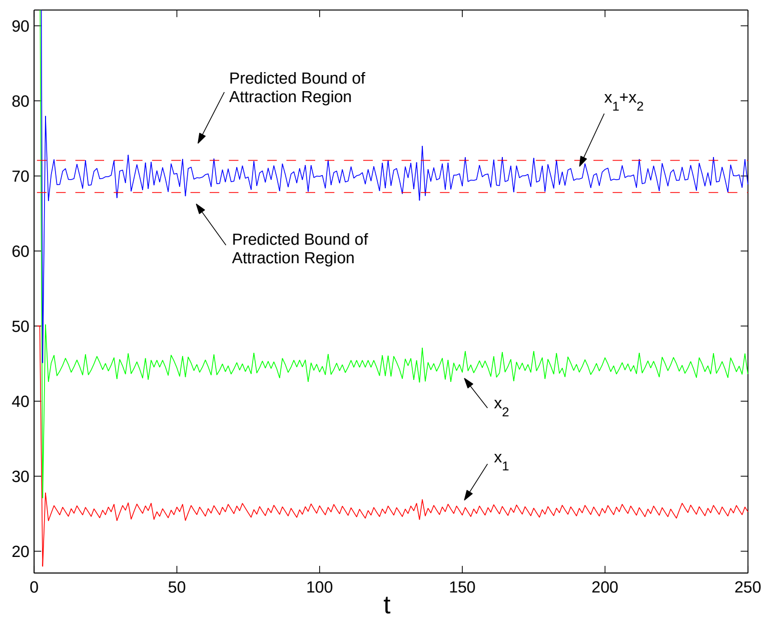

Figure 1.1. Simulation with a two-source one-link network. Note that the sources use different utility functions.

Another common type of error that we will look at is the aforementioned fluctuation error. As in RED or REM, we assume that packets are dropped (or marked) independently, according to the current price (drop probability) on each link. The observed fraction of lost packets at each source is therefore binomial distributed with standard deviation σ

equal to $\frac{\sqrt{q_s(1-q_s)}}{\sqrt{d_s x_s}}$. We make a deterministic 3σ approximation to obtain a bound on $v_s(t)$, and therefore in the single-source-single-link case (1.18) becomes

$$\frac{3\alpha\sqrt{q_s(1-q_s)}}{\sqrt{d_s x_s}} > \eta\,|c - x(t)|$$

This implies

$$\frac{3\alpha}{2}\frac{1}{\sqrt{d_s x_s}} > \eta\,|c - x(t)|$$

since $\sqrt{q_s(1-q_s)} \leq \frac{1}{2}$. By squaring both sides we obtain a cubic function of x_s, and the solution to the inequality can be readily computed. Figure 1.2 shows a simulation result in a single-source-single-link network where fluctuation error is present.

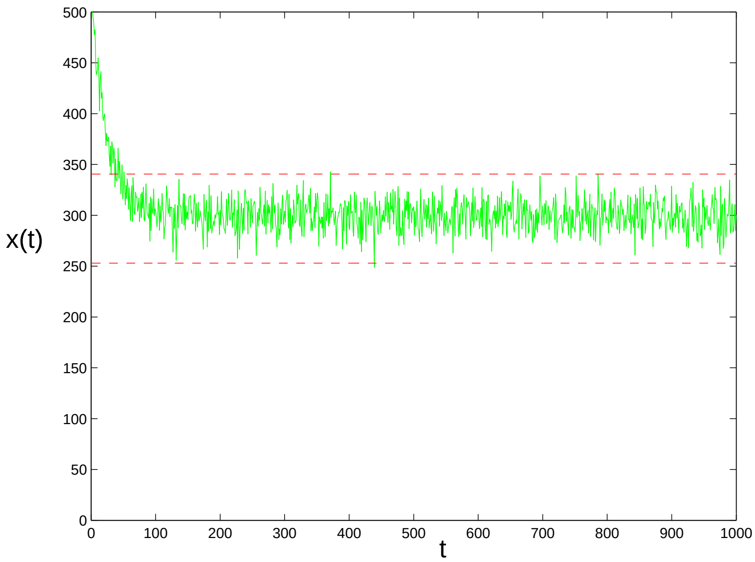

Figure 1.2. Aggregate rate dynamics with probabilistic packet dropping and the 3σ attraction region.

Finally, we present a single-source-single-link simulation illustrating the stepsize dependence discussed in Theorem 1.4. Here we again use the quantization error from our first simulation, and show the rate dynamics for various choices of the stepsize. We indeed observe that all choices below γ_{opt} are outperformed. We do not see any clear superiority to larger values, but we again remark that the optimality is based on worst-case analysis. Finally, we note that increasingly large oscillations ensue as we increase the stepsize. Although this behavior is typical of fixed-stepsize iterative schemes, it is not guaranteed by our analysis.

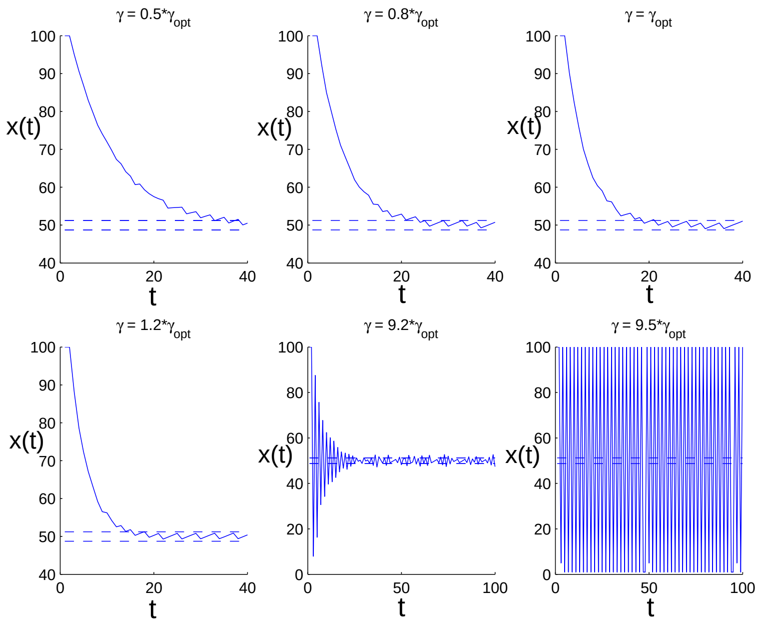

Figure 1.3. Convergence behavior with different choices of γ when quantization error is present.

5. Summary

In this article, we review dual-based modeling and algorithm design in congestion control applications, and study the effects of error in price estimation, which inevitably arises in real networks.

A wide class of source and link algorithms naturally give rise to a network optimization problem through their equilibrium properties. This provides a useful mathematical formalism for interpreting the behavior of various congestion control algorithms, but relies on the assumption that some equilibrium exists. The dual problem, through its distributed structure, naturally inspires source and link algorithms which provably converge to an equilibrium at the optimal point.

These algorithms, however, assume that sources are explicitly notified of the precise link prices. We have examined the effects of imperfect price communication in the dual-gradient algorithm. We utilized the fact that price estimation error was equivalent to an inexact gradient calculation, and hence were able to characterize the dynamics of the network in terms of a region containing the optimum.

Acknowledgments

We acknowledge the financial support of NSF, ARO, and AFOSR.

References

[1] Sanjeewa Athuraliya, Victor H. Li, Steven H. Low, and Qinghe Yin. REM: active queue management. *IEEE Network*, 15(3):48–53, May/June 2001. Extended version in *Proceedings of ITC17*, Salvador, Brazil, September 2001. `http://netlab.caltech.edu`.

[2] D. Bertsekas. *Nonlinear Programming.* Athena Scientific, 1995.

[3] Lawrence S. Brakmo and Larry L. Peterson. TCP Vegas: end-to-end congestion avoidance on a global Internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–80, October 1995. `http://cs.princeton.edu/nsg/papers/jsac-vegas.ps`.

[4] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and distributed computation.* Prentice-Hall, 1989.

[5] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Trans. on Networking*, 1(4):397–413, August 1993. `ftp://ftp.ee.lbl.gov/papers/early.ps.gz`.

[6] V. Jacobson. Congestion avoidance and control. *Proceedings of SIGCOMM'88, ACM*, August 1988. An updated version is available via `ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z`.

[7] Frank P. Kelly, Aman Maulloo, and David Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of Operations Research Society*, 49(3):237–252, March 1998.

[8] Steven H. Low and David E. Lapsley. Optimization flow control, I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, December 1999. `http://netlab.caltech.edu`.

[9] Steven H. Low. A duality model of TCP and queue management algorithms. In *IEEE/ACM Trans. on Networking, 11(4):525-536*, August 2003.

[10] Steven H. Low, Larry Peterson, and Limin Wang. Understanding Vegas: a duality model. *J. of ACM*, 49(2):207–235, March 2002. http://netlab.caltech.edu.

[11] David G. Luenberger. *Linear and Nonlinear Programming, 2nd Ed.* Addison-Wesley Publishing Company, 1984.

[12] Mortada Mehyar, Demetri Spanos, and Steven H. Low. Optimization Flow Control with Estimation Error. In *Infocom'04*, March 2004. http://netlab.caltech.edu/pub/papers/estimation-infocom04.pdf.

[13] Fernando Paganini, John C. Doyle, and Steven H. Low. Scalable laws for stable network congestion control. In *Proceedings of Conference on Decision and Control*, December 2001. http://www.ee.ucla.edu/~paganini.

[14] W. Stevens. *TCP/IP illustrated: the protocols*, volume 1. Addison–Wesley, 1999. 15th printing.