Optimization and Stability of TCP/IP with Delay-Sensitive Utility Functions

Thesis by

John Pongsajapan

In Partial Fulfillment of the Requirements
for the Degree of
Master of Science

California Institute of Technology
Pasadena, California

2006
(Submitted June 2, 2006.)

Acknowledgements

I thank my advisor, Professor Steven Low, for his extraordinary kindness, support, and advice throughout the years.

I am grateful to my family for their love and encouragement.

Abstract

TCP/IP can be interpreted as a distributed primal-dual algorithm to maximize aggregate utility over source rates. It has recently been shown that an equilibrium of TCP/IP, if it exists, maximizes the same aggregate utility function over both source rates and routes, provided pure congestion prices are used as link costs in the shortest-path calculation of IP. Moreover, the utility functions are delay-insensitive, i.e., they are functions of rates only. We extend this result in several ways.

First, we show that if utility functions are delay-insensitive, then there are networks for which TCP/IP optimizes aggregate utility only if routing is based on pure congestion prices. Routing based on the weighted sum $ap_l + \tau_l$ of congestion prices p_l and propagation delays τ_l optimizes aggregate utility for general networks only if the utility functions are not delay-insensitive. Moreover, we identify a class of delay-sensitive utility functions that is implicitly optimized by TCP/IP. As for the delay-insensitive case, we show for this class of utility functions, equilibrium of TCP/IP exists if and only if the optimization problem has zero duality gap. In that case, there is no penalty for not splitting the traffic. We exhibit some counter-intuitive properties of this class of utility functions. We also prove that any class of delay-sensitive utility functions that are optimized by TCP/IP necessarily possess some strange properties.

We prove that, for general networks, if the weight a is small enough, only minimum-propagation-delay paths are selected. Hence if all source-destination pairs have unique minimum-propagation-delay paths, then equilibrium of TCP/IP exists and is asymptotically stable. For general networks, their equilibrium properties are the same as a modified network where paths with non-minimum propagation delays are deleted and routing is based on pure congestion prices.

It is commonly believed that there is generally an inevitable tradeoff between utility maximization and stability in TCP/IP networks. In particular, as the weight a increases, the routing will change from stable to unstable. We exhibit a counterexample where routing changes from stable to unstable and then to stable again, as the weight a increases. Moreover, one can construct a network with any given utility profile as a function of the weight a.

Contents

Chapter 1

Model and Overview

We use the same model as in [5]. In general, we use small letters to denote vectors, e.g., x with x_i as its ith component; capital letters to denote matrices, e.g., H, W, R, or constants, e.g., L, N, K^i; and script letters to denote sets of vectors or matrices, e.g., $\mathcal{W}_s$, $\mathcal{W}_m$, $\mathcal{R}_s$, $\mathcal{R}_m$. Superscript is used to denote vectors, matrices, or constants pertaining to source i, e.g., y^i, w^i, H^i, K^i.

1.1 Network

A network is modeled as a set of L uni-directional links shared by a set of N source-destination pairs, indexed by i (we will also refer to the pair simply as "source i"). Each link l has a finite capacity $c_l > 0$ and a delay $\tau_l > 0$ across the link, i.e., it takes τ_l to process and propagate a packet from one end of the link to the other, excluding queueing delay. Let $c = (c_l, l = 1, \ldots, L)$ and $\tau = (\tau_l, l = 1, \ldots, L)$.

There are K^i acyclic paths for source i represented by a $L \times K^i$ 0-1 matrix H^i where

$$H^i_{lj} = \begin{cases} 1, & \text{if path } j \text{ of source } i \text{ uses link } l \\ 0, & \text{otherwise.} \end{cases}$$

Let $\mathcal{H}^i$ be the set of all columns of H^i that represents all the available paths to i under single-path routing. Define the $L \times K$ matrix H as

$$H = [H^1 \ \ldots \ H^N]$$

where $K := \sum_i K^i$. H defines the topology of the network.

Let w^i be a $K^i \times 1$ vector where the jth entry represents the fraction of i's flow on its jth path such that

$$w^i_j \geq 0 \ \ \forall j, \quad \text{and} \quad \mathbf{1}^T w^i = 1,$$

where $\mathbf{1}$ is a vector of an appropriate dimension with the value 1 in every entry. We require $w_j^i \in \{0,1\}$ for single-path routing, and allow $w_j^i \in [0,1]$ for multi-path routing. Collect the vectors w^i, $i = 1, \ldots, N$, into a $K \times N$ block-diagonal matrix W. Let $\mathcal{W}_s$ be the set of all such matrices corresponding to single path routing defined as

$$\{W | W = \mathrm{diag}(w^1, \ldots, w^N) \in \{0,1\}^{K \times N}, \mathbf{1}^T w^i = 1 \}.$$

Define the corresponding set $\mathcal{W}_m$ for multi-path routing as

$$\{W| \ W = \mathrm{diag}(w^1, \ldots, w^N) \in [0,1]^{K \times N}, \ \mathbf{1}^T w^i = 1 \}. \tag{1.1}$$

As mentioned above, H defines the set of acyclic paths available to each source, and represents the network topology. W defines how the sources load balance across these paths. Their product defines a $L \times N$ routing matrix $R = HW$ that specifies the fraction of i's flow at each link l. The set of all single-path routing matrices is

$$\mathcal{R}_s \quad = \quad \{ \ R \mid R = HW, W \in \mathcal{W}_s \ \}, \tag{1.2}$$

and the set of all multi-path routing matrices is

$$\mathcal{R}_m \quad = \quad \{ \ R \mid R = HW, W \in \mathcal{W}_m \ \}. \tag{1.3}$$

The difference between single-path routing and multi-path routing is the integer constraint on W and R. A single-path routing matrix in $\mathcal{R}_s$ is an 0-1 matrix:

$$R_{li} \quad = \quad \begin{cases} 1, & \text{if link } \ l \text{ is in a path of source } i \\ 0, & \text{otherwise.} \end{cases}$$

A multi-path routing matrix in $\mathcal{R}_m$ is one whose entries are in the range $[0,1]$:

$$R_{li} \begin{cases} > \ 0, & \text{if link } \ l \text{ is in a path of source } i \\ = \ 0, & \text{otherwise.} \end{cases}$$

The path of source i is denoted by $r^i = [R_{1i} \quad \ldots \quad R_{Li}]^T$, the ith column of the routing matrix R.

1.2 TCP–AQM/IP

We consider the situation where TCP–AQM operates at a faster timescale than routing updates. We assume a *single* path is selected for each source-destination pair that minimizes the sum of the link costs in the path, for some appropriate definition of link cost. In particular, traffic is not split across multiple paths from the source to the destination even if they are available. This models, e.g., IP routing within an Autonomous System. We focus on the timescale of the route changes, and assume TCP–AQM is stable and converges instantly to equilibrium after a route change. As in [3], we will interpret the equilibria of various TCP and AQM algorithms as solutions of a utility maximization problem defined in [2]. Different TCP algorithms solve the same prototypical problem (1.4) with different utility functions [3, 4].

Specifically, suppose each source i has a utility function $U_i(x_i, d_i)$ which depends on both its (total transmission) rate x_i and the end-to-end propagation delay d_i. Here, we assume that $d_i = \sum_l R_{li}\tau_l$ and hence the delay d_i depends only on routing R. Given R, $U_i(x_i, d_i)$ is a function only of rate x_i. We assume that utility functions are strictly concave for fixed d_i. We will also consider the special case where the utility function $U_i(x_i) = U_i(x_i, d_i)$ depends on its rate x_i but not on the delay d_i. We will call the first type of utility functions $U_i(x_i, d_i)$ *delay-sensitive* and the second type $U_i(x_i)$ *delay-insensitive*. Delay-insensitive utility functions are studied in Chapter 2 and in [5]; delay-sensitive utility functions are studied in Chapter 3. Throughout this paper, we assume that a network's sources either all have delay-insensitive utility functions, or all have delay-sensitive utility functions.

Let $R(t) \in \mathcal{R}_s$ be the (single-path) routing in period t. Given a $R(t)$, let the equilibrium rates $x(t) = x(R(t))$ and prices $p(t) = p(R(t))$ generated by TCP–AQM in period t, respectively, be the optimal solutions of the constrained maximization problem

$$\max_{x\geq 0} \quad \sum_i U_i(x_i, d_i) \qquad \text{s. t.} \quad R(t)x \ \leq \ c, \tag{1.4}$$

and its Lagrangian dual

$$\min_{p\geq 0} \sum_i \max_{x_i\geq 0} \left(U_i(x_i, d_i) - x_i \sum_l R_{li}(t) p_l \right) \ + \ \sum_l c_l p_l \tag{1.5}$$

The prices $p_l(t)$, $l = 1, \ldots, L$, are measures of congestion, such as queueing delays or loss probabilities [3, 4]. We assume the link costs in period t are

$$z_l(t) \quad = \quad a p_l(t) + b\tau_l \tag{1.6}$$

where $a \geq 0$, $b \geq 0$, and $\tau_l > 0$ are constants. Based on these costs, each source computes its new route $r^i(t+1) \in \mathcal{H}^i$ individually that minimizes the sum of link cost in its path:

$$r^i(t+1) \quad = \quad \arg\min_{r^i\in\mathcal{H}^i} \sum_l z_l(t) r_l^i. \tag{1.7}$$

In (1.6), τ_l are propagation delays across links l. If $p_l(t)$ represents the queueing delays at links l and $a_l = b_l = 1$, then $z_l(t)$ represent total delays across links l. The protocol parameters a and b determine the responsiveness of routing to network traffic: $a = 0$ corresponds to static routing, $b = 0$ corresponds to purely dynamic routing, and the larger the ratio of a/b, the more responsive routing is to network traffic. They determine whether an equilibrium exists, whether it is stable, and the achievable utility at equilibrium. The paper [5] focuses on the case of $b = 0$; we study the general case here.

An equivalent way to specify the TCP–AQM/IP system as a dynamical system, at the timescale of route changes, is to replace (1.4)–(1.5) by their optimality conditions. The routing is updated according to

$$r^i(t+1) = \arg \min_{r^i \in \mathcal{H}^i} \sum_l (a p_l(t) + b\tau_l) r^i_l, \quad \text{for all } i \tag{1.8}$$

where $p(t)$ and $x(t)$ are given by

$$\sum_l R_{li}(t) p_l(t) \quad = \quad \left[\frac{\partial U_i}{\partial x_i}(x_i(t), d_i)\right]^+ \qquad \text{for all } i \tag{1.9}$$

$$\sum_i R_{li}(t) x_i(t) \begin{cases} \le c_l & \text{if } p_l(t) \ge 0 \\ = c_l & \text{if } p_l(t) > 0 \end{cases} \quad \text{for all } l \tag{1.10}$$

$$x(t) \ge 0, \quad p(t) \ge 0. \tag{1.11}$$

This set of equations describe how the routing $R(t)$, rates $x(t)$, and prices $p(t)$ evolve. Note that $x(t)$ and $p(t)$ depend only on $R(t)$ through (1.9)–(1.11), implicitly assuming that TCP–AQM converges instantly to an equilibrium given the new routing $R(t)$.

We say that (R^*, x^*, p^*) is an *equilibrium of TCP/IP* if it is a fixed point of (1.4)–(1.7), or equivalently, (1.8)–(1.11), i.e., starting from routing R^* and associated (x^*, p^*), the above iterations yield (R^*, x^*, p^*) in the subsequent periods.

1.3 Overview of results

TCP/IP can be interpreted as a distributed primal-dual algorithm to maximize aggregate utility over source rates. It has recently been shown that an equilibrium of TCP/IP, if it exists, maximizes the same aggregate utility function over both source rates and routes, provided pure congestion prices are used as link costs in the shortest-path calculation of IP. Moreover, the utility functions are delay-insensitive, i.e., they are functions of rates only. We extend this result in several ways.

First, we show that if utility functions are delay-insensitive, then there are networks for which TCP/IP optimizes aggregate utility only if routing is based on pure congestion prices. Routing based on the weighted sum $ap_l + \tau_l$ of congestion prices p_l and propagation delays τ_l optimizes aggregate utility for general networks only if the utility functions are not delay-insensitive. Moreover, we identify a class of delay-sensitive utility functions that is implicitly optimized by TCP/IP. As for the delay-insensitive case, we show for this class of utility functions, equilibrium of TCP/IP exists if and only if the optimization problem has zero duality gap. In that case, there is no penalty for not splitting the traffic. We exhibit some counter-intuitive properties of this class of utility functions. We also prove that any class of delay-sensitive utility functions that are optimized by TCP/IP necessarily possess some strange properties.

We prove that, for general networks, if the weight a is small enough, only minimum-propagation-delay paths are selected. Hence if all source-destination pairs have unique minimum-propagation-delay paths, then equilibrium of TCP/IP exists and is asymptotically stable. For general networks, their equilibrium properties are the same as a modified network where paths with non-minimum propagation delays are deleted and routing is based on pure congestion prices.

It is commonly believed that there is generally an inevitable tradeoff between utility maximization and stability in TCP/IP networks. In particular, as the weight a increases, the routing will change from stable to unstable. We exhibit a counterexample where routing changes from stable to unstable and then to stable again, as the weight a increases. Moreover, one can construct a network with any given utility profile as a function of the weight a.

Chapter 2

Delay-Insensitive Network Optimization

In this chapter, we consider the special case where the utility functions $U_i(x_i)$ depend only on rates x_i and not on propagation delays d_i. It is shown in [5] that TCP/IP maximizes aggregate utility over both rates and routing when an equilibrium exists, provided that pure price is used as the link cost in shortest-path routing, i.e., $b = 0$ in (1.8). We now argue that in the other case ($b > 0$ for all l in (1.8)), TCP/IP in general does not maximize aggregate utility. In the next chapter, we will show that TCP/IP turns out to maximize a class of delay-sensitive utility functions when $b > 0$.

2.1 The optimization problem

Definition 1 *A* delay-insensitive *utility function is a strictly concave increasing, continuously differentiable function $U_i(x_i)$ from $[0, \infty)$ to $[-\infty, \infty]$.*

In this chapter, we restrict ourselves to delay-insensitive utility functions.

Consider the single-path network optimization problem from [5]:

$$\max_{R \in \mathcal{R}_s, x \geq 0} \sum_{i=1}^{N} U_i(x_i) \quad \text{s.t. } Rx \leq c \tag{2.1}$$

and its Lagrangian dual:

$$\min_{p \geq 0} \sum_{i=1}^{N} \max_{x_i \geq 0} \left(U_i(x_i) - x_i \min_{r^i \in \mathcal{H}^i} \sum_{l=1}^{L} R_{li} p_l \right) + \sum_{l=1}^{L} c_l p_l \tag{2.2}$$

where r^i is the ith column of R with $r^i_l = R_{li}$. This problem maximizes utility over both rates and routes.

Define the Lagrangian [1]:

$$L(R, x, p) = \sum_{i=1}^{N} \left(U_i(x_i) - x_i \min_{r^i \in \mathcal{H}^i} \sum_{l=1}^{L} R_{li} p_l \right) + \sum_{l=1}^{L} c_l p_l$$

Then we can express the primal and dual problems respectively as:

$$V_{sp} = \max_{R \in \mathcal{R}_s} \max_{x \geq 0} \min_{p \geq 0} L(R, x, p)$$

$$V_{sd} = \min_{p \geq 0} \max_{R \in \mathcal{R}_s} \max_{x \geq 0} L(R, x, p)$$

If we allow sources to use multiple paths, the corresponding problems are:

$$V_{mp} = \max_{R \in \mathcal{R}_s} \max_{x \geq 0} \min_{p \geq 0} L(R, x, p)$$

$$V_{md} = \min_{p \geq 0} \max_{R \in \mathcal{R}_s} \max_{x \geq 0} L(R, x, p)$$

The TCP/IP dynamical system is described by (1.4)–(1.7), or equivalently, (1.8)–(1.11) with $U_i(x_i, d_i)$ replaced by $U_i(x_i)$ and $\partial U_i / \partial x_i$ by $U_i'(x_i)$.

2.2 Optimization with link cost p_l $(b = 0)$

In this section, we assume that IP uses the congestion prices p_l generated by TCP-AQM as link costs, i.e., $a > 0$ and $b = 0$ in (1.8).

From [5], we know the following:

Theorem 1 *Suppose $a > 0$ and $b = 0$ in (1.8). Then:*

1. *An equilibrium (R^*, x^*, p^*) of TCP/IP exists if and only if there is no duality gap between* (2.1) *and* (2.2).

2. *In this case, the equilibrium (R^*, x^*, p^*) is a solution of* (2.1) *and* (2.2).

Moreover,

Theorem 2

$$V_{sp} \leq V_{sd} = V_{mp} = V_{md}.$$

2.3 Optimization with link cost τ_l $(a = 0)$

In this section, we assume that IP uses the propagation delays τ_l as link costs, i.e., $a = 0$ and $b > 0$ in (1.8).

Theorem 3 *Suppose $a = 0$ and $b > 0$ in (1.8). Then for any delay-insensitive utility function $U(x)$, there exists a network with sources using this utility function, where TCP/IP equilibrium exists but does not solve* (2.1) *and* (2.2).

Proof. Suppose $c_2 > c_1 > 0$ and $\tau_2 > \tau_1 > 0$. Consider the following network.

Network 1:

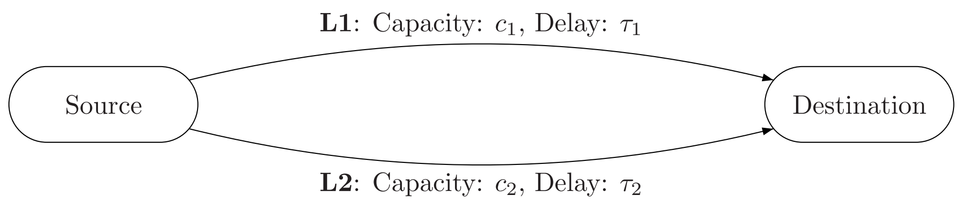

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x)$

Flow 1 achieves rate c_1 with utility $U(c_1)$.

The route costs are:

Route Costs	
Route	Cost
R1	$b\tau_1$
R2	$b\tau_2$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 1 had instead used route **R2**. Then it would achieve rate c_2, with utility $U(c_2)$. But $c_1 < c_2$, so $U(c_1) < U(c_2)$.

So there exists a TCP/IP equilibrium that does not solve (2.1) and (2.2). $\square$

2.4 Optimization with link cost $ap_l + \tau_l$

In this section, we assume that IP uses $ap_l + \tau_l$ as link cost, i.e., $a > 0$ and $b = 1$ in (1.8).

Theorem 4 *Suppose $a > 0$ and $b = 0$ in (1.8). Then for any delay-insensitive utility function $U(x)$, there exists a network with sources using this utility function, where TCP/IP equilibrium exists but does not solve* (2.1) *and* (2.2).

Proof. Suppose $c_2 > c_1 > 0$, and $\tau > 0$. Consider the following network.

Network 2:

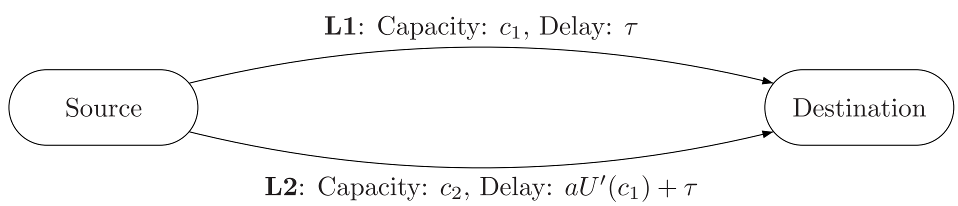

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x)$

Flow 1 achieves rate c_1 with utility $U(c_1)$.

The route costs are then:

Route Costs	
Route	Cost
R1	$aU'(c_1) + \tau$
R2	$aU'(c_1) + \tau$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 1 had instead used route **R2**. Then it would achieve rate c_2, with utility $U(c_2)$. But $c_1 < c_2$, so $U(c_1) < U(c_2)$.

So there exists a TCP/IP equilibrium that does not solve (2.1) and (2.2). $\square$

Chapter 3

Delay-Sensitive Network Optimization

In this chapter, we consider the case where the utility functions $U_i(x_i, d_i)$ depend only both on rates x_i and on propagation delays d_i. We identify a class $\mathcal{C}$ of utility functions for which TCP/IP, with $a > 0$ and $b = 1$ in (1.8), does maximize aggregate utility at equilibrium, when equilibrium exists. We argue that for the other cases, $a = 0$ or $b = 0$, TCP/IP in general does not maximize aggregate utility at equilibrium. We analyze properties of $\mathcal{C}$, and then derive properties that any class of utility functions that TCP/IP implicitly maximizes at equilibrium must possess.

3.1 The optimization problem

Definition 2 *A* delay-sensitive *utility function is a continuously differentiable function $U(x, d)$ from $[0, \infty) \times [0, \infty)$ to $[-\infty, \infty)$, that satisfies the following properties:*

1. *$\forall$ fixed $d > 0$, $U(x, d)$ is strictly concave in x.*
2. *$\forall d > 0, x > 0$, $U(x, d)$ and $\frac{\partial U}{\partial x}(x, d)$ are finite.*
3. *$(\forall x > 0), (\forall d_1, d_2 \text{ s.t. } 0 < d_1 < d_2) : U(x, d_1) > U(x, d_2)$*
4. *$(\exists D > 0), (\forall 0 < d < D), (\exists X(d) > 0), (\forall x < X(d)) : \frac{\partial U}{\partial x}(x, d) > 0$*

Essentially, a delay-sensitive utility function is defined so that the source always gains utility from reducing propagation delay. If propagation delay is too high, the source can choose not to transmit. Otherwise, for fixed delay, the source's utility increases with transmission rate, possibly up to some limit. This may seem like an unusually permissive definition, but we will show that it provides some useful results.

Theorem 5 *For every delay-sensitive function $U(x, d)$, $(\exists D > 0), (\forall 0 < d < D), (\exists X(d) > 0)$, $(\forall x_1, x_2 \text{ s.t. } 0 < x_1 < x_2 < X(d))$: $U(x_1, d) < U(x_2, d)$ and $\frac{\partial U}{\partial x}(x_1, d) > \frac{\partial U}{\partial x}(x_2, d) > 0$.*

Proof. The definition of delay-sensitive function guarantees that:

$(\exists D > 0), (\forall 0 < d < D), (\exists X(d) > 0) : \frac{\partial U}{\partial x}(X(d), d) > 0.$

But then the rest follows since $U(x, d)$ is continuously differentiable and strictly concave for fixed d. □

We assume all sources on the network have delay-sensitive utility functions $U_i(x_i, d_i)$ where x_i is its transmission rate and d_i is its path delay, given by

$$d_i = \sum_{l=1}^{L} R_{li}\tau_l,$$

where the routing matrix R is in $\mathcal{R}_s$ for single-path routing and in $\mathcal{R}_m$ for multi-path routing. Note that in the multi-path case, d_i is the traffic weighted average of propagation delays along its paths.

We adapt the single-path delay-insensitive network optimization problem from [5] to a delay-sensitive network optimization problem:

$$\max_{R\in\mathcal{R}_s, x\geq 0} \sum_{i=1}^{N} U_i\left(x_i, \sum_{l=1}^{L} R_{li}\tau_l\right) \quad \text{s.t. } Rx \leq c, \tag{3.1}$$

Its Lagrangian dual is:

$$\min_{p\geq 0} \sum_{i=1}^{N} \max_{x_i\geq 0} \max_{r^i\in\mathcal{H}^i} \left(U_i(x_i, d_i) - x_i \sum_{l=1}^{L} R_{li}p_l\right) + \sum_{l=1}^{L} c_l p_l, \tag{3.2}$$

where r^i is the ith column of R with $r^i_l = R_{li}$. This problem maximizes utility over both rates and routes.

Define the Lagrangian [1]:

$$L(R, x, p) = \sum_{i=1}^{N} \left(U_i(x_i, d_i) - x_i \sum_{l=1}^{L} R_{li}p_l\right) + \sum_{l=1}^{L} c_l p_l.$$

Then we can express the primal and dual problems respectively as:

$$V_{sp} = \max_{R\in\mathcal{R}_s} \max_{x\geq 0} \min_{p\geq 0} L(R, x, p)$$

$$V_{sd} = \min_{p\geq 0} \max_{R\in\mathcal{R}_s} \max_{x\geq 0} L(R, x, p)$$

If we allow sources to use multiple paths, the corresponding problems are:

$$V_{mp} = \max_{R \in \mathcal{R}_s} \max_{x \geq 0} \min_{p \geq 0} L(R, x, p)$$

$$V_{md} = \min_{p \geq 0} \max_{R \in \mathcal{R}_s} \max_{x \geq 0} L(R, x, p)$$

The TCP/IP dynamical system is described by (1.4)–(1.7), or equivalently, (1.8)–(1.11).

3.2 Optimization with link cost p_l ($b = 0$)

In this section, we assume that IP uses the congestion prices p_l generated by TCP-AQM as link costs, i.e., $a > 0$ and $b = 0$ in (1.8).

Theorem 6 *Suppose $a > 0$ and $b = 0$ in (1.8). Then for every delay-sensitive utility function $U(x, d)$, there exists a network with sources using this utility function, where TCP/IP equilibrium exists but does not solve* (3.1) *and* (3.2).

Proof. Suppose we have a delay-sensitive utility function $U(x, d)$. From Theorem 5, there exist $\tau_2 > \tau_1 > 0$, $c > 0$ so that $\frac{\partial U}{\partial x}(c, \tau_1) > 0$ and $\frac{\partial U}{\partial x}(c, \tau_2) > 0$. Consider the following network.

Network 3:

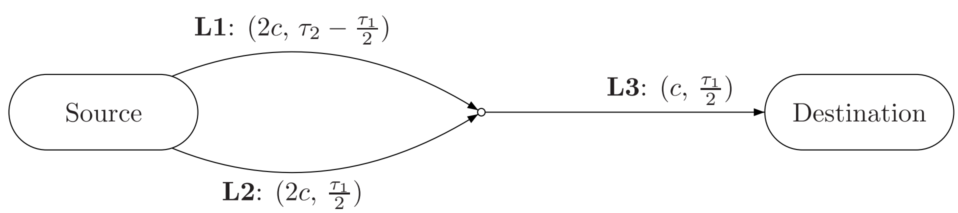

Network Routes	
Route	Path
R1	**L1**, **L3**
R2	**L2**, **L3**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x, d)$

Flow 1 is constrained by **L3** and achieves rate c with propagation delay τ_2 and utility $U(c, \tau_2)$. The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(c, \tau_2)$
R2	$a\frac{\partial U}{\partial x}(c, \tau_2)$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 1 had instead used route **R2**. Then it would achieve rate c with propagation delay τ_1 and utility $U(c, \tau_1)$. But $U(c, \tau_2) < U(c, \tau_1)$.

So there exists a TCP/IP equilibrium that does not solve (3.1) and (3.2). □

3.3 Optimization with link cost τ_l $(a = 0)$

In this section, we assume that IP uses the propagation delays τ_l as link costs, i.e., $a = 0$ and $b > 0$ in (1.8).

Theorem 7 *Suppose $a = 0$ and $b > 0$ in (1.8). Then for any delay-sensitive utility function $U(x, d)$, there exists a network with sources using this utility function, where TCP/IP equilibrium exists but does not solve* (3.1) *and* (3.2).

Proof. Suppose we have a delay-sensitive utility function $U(x, d)$. From Theorem 5, there exist $\tau > 0, c_2 > c_1 > 0$ so that:

1. $U(c_1, \tau) < U(c_2, \tau)$
2. $\frac{\partial U}{\partial x}(c_1, \tau) > \frac{\partial U}{\partial x}(c_2, \tau) > 0$.

But since $U(x, d)$ is continuously differentiable, there exists $\epsilon > 0$ so that:

1. $U(c_1, \tau) < U(c_2, \tau + \epsilon)$
2. $\frac{\partial U}{\partial x}(c_2, \tau + \epsilon) > 0$.

Consider the following network.

Network 4:

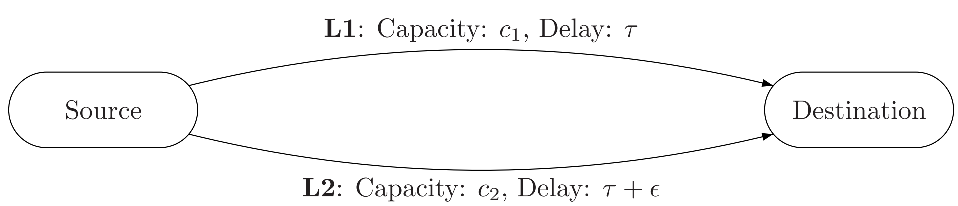

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x, d)$

Flow 1 achieves rate c_1 with propagation delay τ and utility $U(c_1, \tau)$. The route costs are:

Route Costs	
Route	Cost
R1	$b\tau$
R2	$b\tau + b\epsilon$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 1 had instead used route **R2**. Then it would achieve rate c_2 with propagation delay $\tau + \epsilon$ and utility $U(c_2, \tau + \epsilon)$. But $U(c_1, \tau) < U(c_2, \tau + \epsilon)$.

So there exists a TCP/IP equilibrium that does not solve (3.1) and (3.2). $\square$

3.4 Optimization with link cost $ap_l + \tau_l$

In this section, we assume that IP uses $ap_l + \tau_l$ as link cost, i.e., $a > 0$ and $b = 1$ in (1.8). We show that there is a class of delay-sensitive utility functions that TCP/IP implicitly maximizes at equilibrium, when equilibrium exists. Moreover, it seems unlikely that other classes have the same property.

3.4.1 A class of delay-sensitive utility functions

Consider the class $\mathcal{C}$ of functions $U(x, d)$ that can be written as:

$$U(x, d) = V(x) - a^{-1}xd$$

where $V(x)$ is a continuously differentiable function from $[0, \infty)$ to $[-\infty, \infty)$ so that $V(x)$ is strictly concave, and $\forall x > 0$, $V(x)$ and $V'(x)$ are finite.

We first show that functions in $\mathcal{C}$ are delay-sensitive utility functions.

Lemma 1 *Every function $U(x, d) \in \mathcal{C}$, for fixed $d > 0$, is strictly concave in x.*

Proof. This is clear. $\square$

Lemma 2 *$\forall d > 0, x > 0$, $U(x, d)$ and $\frac{\partial U}{\partial x}(x, d)$ are finite.*

Proof. This is clear. $\square$

Lemma 3 $(\forall U(x, d) \in \mathcal{C}), (\forall x > 0), (\forall d_1, d_2$ *s.t.* $0 < d_1 < d_2) : U(x, d_1) > U(x, d_2)$.

Proof. Suppose $U(x,d) \in \mathcal{C}$, $x > 0$, $d_2 > d_1 > 0$. Then

$$\begin{aligned} U(x,d_1) - U(x,d_2) &= (V(x) - a^{-1}d_1x) - (V(x) - a^{-1}d_2x) \\ &= a^{-1}d_2x - a^{-1}d_1x \\ &= a^{-1}(d_2 - d_1)x \\ &> 0 \end{aligned}$$

So $U(x,d_1) > U(x,d_2)$ as desired. □

Lemma 4 $(\forall U(x,d) \in \mathcal{C}), (\exists D > 0), (\forall 0 < d < D), (\exists X(d) > 0), (\forall x < X(d)) : \frac{\partial U}{\partial x}(x,d) > 0$

Proof. Suppose $U(x,d) \in \mathcal{C}$. Then $U(x,d) = V(x) - a^{-1}xd$, where $V(x)$ is continuously differentiable and strictly concave increasing. Choose D so that $0 < D < aV'(0)$. This is possible since $V(x)$ is strictly increasing, so that $V'(0) > 0$. Suppose $0 < d < D$. Since $V(x)$ is continuously differentiable, $\exists X > 0$ so that $V'(X) > a^{-1}d$. Suppose $x < X$. Then:

$$\frac{\partial U}{\partial x}(x,d) = V'(x) - a^{-1}d > a^{-1}d - a^{-1}d = 0,$$

where the inequality holds since V is strictly concave increasing, so that $V'(x) > V'(X) > a^{-1}d$. □

Theorem 8 *Every function $U(x,d) \in \mathcal{C}$ is a delay-sensitive utility function.*

Proof. We use the preceding lemmas. □

3.4.2 Initial analysis

Specializing to utility functions in $\mathcal{C}$, the optimization problem (3.1) reduces to:

$$\max_{R \in \mathcal{R}_s, x \geq 0} \sum_{i=1}^{N} [V_i(x_i) - a^{-1}x_i \sum_{l=1}^{L} R_{li}\tau_l] \quad \text{s.t. } Rx \leq c$$

Consider the Lagrangian

$$\begin{aligned} L(R,x,p) &= \sum_{i}^{N} [V_i(x_i) - a^{-1}x_i \sum_{l}^{L} R_{li}\tau_l)] - \sum_{l}^{L} [p_l((\sum_{i}^{N} R_{li}x_i) - c_l)] \\ &= \sum_{i}^{N} [V_i(x_i) - a^{-1}x_i \sum_{l=1}^{L} R_{li}\tau_l - x_i \sum_{l=1}^{L} R_{li}p_l] + \sum_{l=1}^{L} p_l c_l \\ &= \sum_{i}^{N} [V_i(x_i) - x_i \sum_{l=1}^{L} R_{li}(p_l + a^{-1}\tau_l)] + \sum_{l=1}^{L} p_l c_l \end{aligned}$$

and the dual problem

$$
\begin{aligned}
D(p) &= \max_{R\in\mathcal{R}_s, x\geq 0} L(R,x,p)\\
&= \max_{R\in\mathcal{R}_s, x\geq 0} \sum_{i}^{N} [V_i(x_i) - x_i \sum_{l=1}^{L} R_{li}(p_l + a^{-1}\tau_l)] + \sum_{l=1}^{L} p_l c_l\\
&= \max_{x\geq 0} \sum_{i=1}^{N} [V_i(x_i) - x_i \min_{r^i\in\mathcal{H}^i} \sum_{l=1}^{L} R_{li}(p_l + a^{-1}\tau_l)] + \sum_{l=1}^{L} p_l c_l
\end{aligned}
$$

Lemma 5 *R^i minimizes $\sum_{l=1}^{L} R_{li}(p_l + a^{-1}\tau_l)$ if and only if R^i minimizes $\sum_{l=1}^{L} R_{li}(ap_l + \tau_l)$.*

Proof. This is obvious. □

The minimization over R in the dual problem appears to involve minimal-cost routing using $p_l + a^{-1}\tau_l$ as route cost. But this is the same as minimal-cost routing using $ap_l + \tau_l$ as route cost.

This suggests that TCP/IP might solve the optimization problem, with utility functions in $\mathcal{C}$.

3.4.3 TCP/IP solves the cross-layer optimization problem

In this section we prove a result analogous to Theorem 1.

Lemma 6 *If utility functions are in $\mathcal{C}$, then if $(\tilde{R}, \tilde{x}, \tilde{p})$ is an equilibrium of TCP/IP, then $(\tilde{R}, \tilde{x}, \tilde{p})$ solves* (3.1) *and* (3.2).

Proof. Suppose $(\tilde{R}, \tilde{x}, \tilde{p})$ is an equilibrium of TCP/IP, so that:

$$
(a\tilde{p} + \tau)^T \tilde{r}^i = \min_{r^i\in\mathcal{H}^i} (a\tilde{p} + \tau)^T r^i \quad \text{for all } i, \tag{3.3}
$$

$$
(\tilde{p}, \tilde{x}) = \arg\ \min_{p\geq 0} \max_{x\geq 0} \left(\sum_{i}^{N} V(x_i) - (p + a^{-1}\tau)^T \tilde{R}x \right) + p^T c \tag{3.4}
$$

Note that from Lemma 5, Equation 3.3 implies

$$
(\tilde{p} + a^{-1}\tau)^T \tilde{r}^i = \min_{r^i\in\mathcal{H}^i} (\tilde{p} + a^{-1}\tau)^T r^i \quad \text{for all } i.
$$

Consider the dual optimal solution

$$
(p^*, x^*, R^*) = \arg\ \min_{p\geq 0} \max_{x\geq 0} \left(\sum_{i}^{N} V(x_i) - \min_{R\in\mathcal{R}_s} (p + a^{-1}\tau)^T Rx \right) + p^T c
$$

Let:

$$f(p) := \max_{x\geq 0}\left(\sum_i^N V(x_i) - (p + a^{-1}\tau)^T \tilde{R}x\right) + p^T c$$
$$g(p) := \max_{x\geq 0}\left(\sum_i^N V(x_i) - \min_{R\in\mathcal{R}_s}(p + a^{-1}\tau)^T Rx\right) + p^T c$$

Note that $f(\tilde{p}) = \min_{p\geq 0} f(p)$ and $g(p^*) = \min_{p\geq 0} g(p)$. Also, since $\tilde{R} \in \mathcal{R}_s$, we have

$$f(p) \leq g(p) \quad \text{for all } p \geq 0$$

and so

$$f(\tilde{p}) = \min_{p\geq 0} f(p) \leq \min_{p\geq 0} g(p) = g(p^*)$$

But

$$\begin{aligned}
f(\tilde{p}) &= \max_{x\geq 0}\sum_i^N V(x_i) - (\tilde{p} + a^{-1}\tau)^T \tilde{R}x + \tilde{p}^T c \\
&= \max_{x\geq 0}\sum_i^N V(x_i) - \sum_i^N x_i\left((\tilde{p} + a^{-1}\tau)\tilde{r}^i\right) + \tilde{p}^T c \\
&= \max_{x\geq 0}\sum_i^N V(x_i) - \sum_i^N x_i\left(\min_{r^i\in\mathcal{H}^i}(\tilde{p} + a^{-1}\tau)r^i\right) + \tilde{p}^T c \\
&= g(\tilde{p}) \\
&\geq g(p^*)
\end{aligned}$$

So $f(\tilde{p}) = g(p^*) = g(\tilde{p})$, $L(\tilde{R}, \tilde{x}, \tilde{p}) = L(R^*, x^*, p^*)$, and $(\tilde{R}, \tilde{x}, \tilde{p})$ solves the dual problem. Further, $(\tilde{R}, \tilde{x}, \tilde{p})$ is primal feasible, and so solves the primal problem. $\square$

Lemma 7 *If utility functions are in $\mathcal{C}$, then if (R^*, x^*, p^*) solves* (3.1) *and* (3.2) *with zero duality gap, then (R^*, x^*, p^*) is an equilibrium of TCP/IP.*

Proof. Assuming that there is no duality gap and (R^*, x^*, p^*) is an optimal solution, then we want to show that it is also an equilibrium, so that

$$(ap^* + \tau)^T r^{*i} = \min_{r^i\in\mathcal{H}^i}(ap^* + \tau)^T r^i \tag{3.5}$$
$$\begin{aligned}
(p^*, x^*) &= \arg\ \min_{p\geq 0}\max_{x\geq 0} L(R^*, x, p) \\
&= \arg\ \max_{x\geq 0}\min_{p\geq 0} L(R^*, x, p)
\end{aligned} \tag{3.6}$$

(3.5) follows from the Saddle-Point Theorem [1] and Lemma 5. (3.6) holds since (p^*, x^*) solves the dual of the optimization problem with fixed R^*. □

Theorem 9 *If utility functions are in $\mathcal{C}$:*

1. *An equilibrium (R^*, x^*, p^*) of TCP/IP exists if and only if there is no duality gap between* (2.1) *and* (2.2).
2. *In this case, the equilibrium (R^*, x^*, p^*) is a solution of* (2.1) *and* (2.2).

Proof. We use the preceding lemmas. □

3.4.4 Zero duality gap implies no splitting advantage

In this section we prove a result analogous to Theorem 2.

Lemma 8 *If utility functions are in $\mathcal{C}$, $V_{sp} \leq V_{sd}$.*

Proof. This follows from the weak duality theorem [1]. □

Lemma 9 *If utility functions are in $\mathcal{C}$, $V_{md} = V_{sd}$.*

Proof. Here we show that the utility of the dual optimal solution is the same for both multi-path and single-path problems.

$$V_{sd} = \min_{p\geq 0} \max_{R\in\mathcal{R}_s, x\geq 0} \sum_{i=1}^{N} V_i(x_i) - [(p + a^{-1}\tau)^T Rx] + \sum_{l=1}^{L} [p_l c_l]$$

$$V_{sd} = \min_{p\geq 0} \max_{x\geq 0} \sum_{i=1}^{N} V_i(x_i) - \min_{W\in\mathcal{W}_s} [(p + a^{-1}\tau)^T HWx] + \sum_{l=1}^{L} [p_l c_l]$$

$$V_{md} = \min_{p\geq 0} \max_{x\geq 0} \sum_{i=1}^{N} V_i(x_i) - \min_{W\in\mathcal{W}_m} [(p + a^{-1}\tau)^T HWx] + \sum_{l=1}^{L} [p_l c_l]$$

Define $f_s(x, p)$ and $f_m(x, p)$ as:

$$f_s(x, p) := \min_{W\in\mathcal{W}_s} [(p + a^{-1}\tau)^T HWx]$$

$$f_m(x, p) := \min_{W\in\mathcal{W}_m} [(p + a^{-1}\tau)^T HWx]$$

Then:

$$f_m(x,p) = \min_{W \in \mathcal{W}_m} [(p + a^{-1}\tau)^T HWx] \quad \text{s.t. } \mathbf{1}^T w^i = 1,\ 0 \le w^i_j \le 1.$$

This is a linear optimization problem, so there is an optimal W on the boundary (in $\mathcal{W}_s$). So we can conclude $V_{md} = V_{sd}$. □

Lemma 10 *If utility functions are in* $\mathcal{C}$*,* $V_{md} = V_{mp}$*.*

Proof. Here we show that the multi-path problem has no duality gap.

$$V_{mp} = \max_{R \in \mathcal{R}_m, x \ge 0} \sum_{i=1}^{N} V_i(x_i) - a^{-1} x_i \sum_{l=1}^{L} R_{li}\tau_l \quad \text{s.t. } Rx \le c$$

Define the $K_i \times 1$ vector y^i, for each source i, as follows:

$$y^i := x_i w^i$$

The mapping from (x_i, w^i) to y^i is one to one: $x_i = \mathbf{1}^T y^i$, and $w^i = y^i / x_i$.

Define the $K_i \times 1$ path propagation delay vector d^i, for each source i, so that:

$$d^i_j := \sum_{l=1}^{L} H^i_{lj}\tau_l$$

This vector contains the total propagation delays for source i's paths.

Then, we observe:

$$\begin{aligned}
x_i \sum_{l=1}^{L} R_{li}\tau_l &= x_i \sum_{l=1}^{L}\sum_{j=1}^{K_i} H^i_{lj} w^i_j \tau_l \\
&= x_i \sum_{j=1}^{K_i}\sum_{l=1}^{L} H^i_{lj} w^i_j \tau_l \\
&= x_i \sum_{j=1}^{K_i} w^i_j \sum_{l=1}^{L} H^i_{lj}\tau_l \\
&= x_i (w^i \cdot d^i) \\
&= y^i \cdot d^i
\end{aligned}$$

And also:

$$Rx = HWx = Hy$$

Then we can reformulate the optimization problem as this equivalent problem:

$$\max_{y \geq 0} \sum_{i=1}^{N} V_i(\mathbf{1}^T y^i) - a^{-1}(y^i \cdot d^i) \quad \text{s.t. } Hy \leq c$$

This is a convex program with linear constraint and therefore has no duality gap. So we conclude $V_{mp} = V_{md}$. □

Theorem 10 *If utility functions are in $\mathcal{C}$, $V_{sp} \leq V_{sd} = V_{mp} = V_{md}$.*

Proof. We simply use the lemmas from this section. □

3.4.5 Properties of class C

3.4.5.1 Partial utilization

Theorem 11 *Given any utility function in $\mathcal{C}$, there exists a network where TCP/IP underutilizes links.*

Proof. We will prove this by explicit construction.

Consider:

$$U(x,d) = V(x) - a^{-1}xd$$
$$\frac{\partial U}{\partial x}(x,d) = V_i'(x_i) - a^{-1}d$$

If $V'(x) - a^{-1}d = 0$, then x is the rate that maximizes $U(x,d)$ for fixed d, since $U(x,d)$ is strictly concave for fixed d.

Choose any $c > 0$ and set $\tau = aV'(c)$. Note that $\tau > 0$, since $V(x)$ is strictly increasing. Consider the following network.

Network 5:

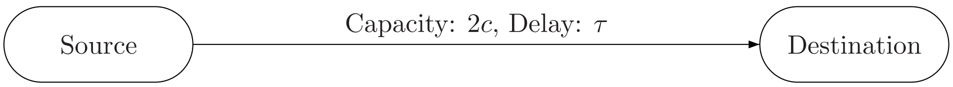

Network Routes	
Route	Path
R1	**L1**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**	**R1**	$U(x, d)$

Evidently, $\frac{\partial U}{\partial x}(x, d) = V'(c) - a^{-1}\tau = V'(c) - V'(c) = 0$, so c is the optimal rate for Flow 1, and therefore the equilibrium rate. But this leaves the link underutilized, since the link has capacity $2c$.

$\square$

3.4.5.2 Non-utilization of extra routes

Consider any $U(x, d) \in \mathcal{C}$, so that $U(x, d) = V(x) - a^{-1}xd$. From Theorem 5, there exist $\tau > 0$, $c > 0$ so that $\frac{\partial U}{\partial x}(c, \tau) > 0$. Consider the following network.

Network 6:

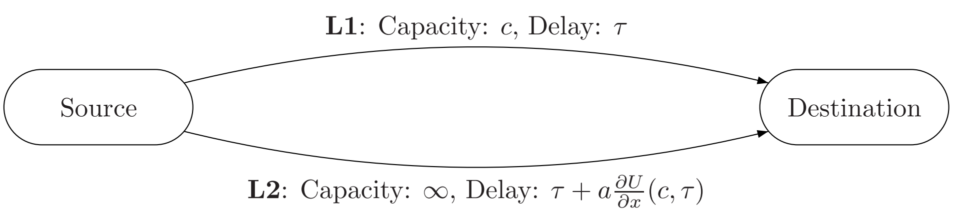

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x, d)$

Flow 1 achieves rate c with propagation delay τ and utility $U(c, \tau)$. The route costs are then:

Route Costs	
Route	Cost
R1	$\tau + a\frac{\partial U}{\partial x}(c,\tau)$
R2	$\tau + a\frac{\partial U}{\partial x}(c,\tau)$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Theorem 9 and Theorem 10 imply that there is no duality gap, and therefore there is no benefit in multi-path routing, and therefore no benefit in utilizing route **R2**. This seems counterintuitive – with all delay-insensitive networks, there is always a benefit in utilizing previously unutilized routes. However, utilizing an extra route can increase the average propagation delay experienced by the flow, which, depending on the utility function, could be suboptimal, regardless of the amount of extra throughput.

We can show this directly.

Remark 1 *Assuming utility functions in $\mathcal{C}$, it is suboptimal to use route* ***R2****, even when the flow is allowed to distribute its traffic over multiple paths.*

Proof. Let kc specify the throughput on link **L2**, so that $\frac{k}{k+1}$ specifies the fraction of traffic sent over link **L2** (where the rest is sent over link **L1**). Then the total throughput is given by $kc+c$, and the weighted average propagation delay is:

$$\frac{k\left[\tau + a\frac{\partial U}{\partial x}(c,\tau)\right] + \tau}{k+1}.$$

So for every $k > 0$, we must have:

$$U(c,\tau) \geq U\left(kc + c, \frac{k\left[\tau + a\frac{\partial U}{\partial x}(c,\tau)\right] + \tau}{k+1}\right)$$

Plugging in our utility function yields:

$$
\begin{aligned}
V(x) - a^{-1}c\tau &\geq U\left(kc + c, \frac{k\left(\tau + a\frac{\partial U}{\partial x}(c,\tau)\right) + \tau}{k+1}\right) \\
&\geq V(kc+c) - a^{-1}\left(\frac{(kc+c)\left(k\left(a(V'(c) - a^{-1}\tau) + \tau\right) + \tau\right)}{k+1}\right) \\
&\geq V(kc+c) - a^{-1}\left(\frac{(k+1)c\left(kaV'(c) + \tau\right)}{k+1}\right) \\
&\geq V(kc+c) - a^{-1}c\left(kaV'(c) + \tau\right) \\
&\geq V(kc+x) - kcV'(x) - a^{-1}c\tau \\
V(c) &\geq V(kc+c) - kcV'(c) \\
V'(c) &\geq \frac{V(kc+c) - V(c)}{kc}
\end{aligned}
$$

Substituting $y = kc$, we end up with:

$$
V'(c) \geq \frac{V(c+y) - V(c)}{y}
$$

But this is true since $V(x)$ is concave. $\square$

3.4.6 Alternative classes

In this section, we analyze properties of classes of delay-sensitive utility functions that TCP/IP implicitly optimizes at equilibrium, using $ap_l + \tau_l$ as link cost. In particular, we do not restrict our analysis to class $\mathcal{C}$.

Class $\mathcal{C}$ seems unusual in some respects. Here are two of them:

1. $\exists U(x,d) \in \mathcal{C}, d > 0$ so that $U(x,d)$ is not strictly increasing in x.

2. $\forall U_1(x,d) \in \mathcal{C}$, $\epsilon > 0$: $U_2 := U_1(x+\epsilon)$ is not in $\mathcal{C}$.

We will show that any class of delay-sensitive utility functions that TCP/IP implicitly maximizes at equilibrium cannot avoid both of these unusual properties without introducing another unusual property.

Throughout this section, we use the following function:

$$
M(U,d) := \lim_{c\to\infty} U(c,d), \tag{3.7}
$$

which computes the maximum possible utility at each delay d for a delay-sensitive utility function $U(x,d)$.

Lemma 11 *Suppose that for all fixed* $d > 0$*, a delay-sensitive utility function* $U(x, d)$ *is strictly increasing in* x*. Further suppose that on all networks with sources using this utility function, if TCP/IP equilibrium exists, the equilibrium solves* (3.1) *and* (3.2)*. Then* $\forall x > 0, d > 0 : U(x, d) \geq M(U, a\frac{\partial U}{\partial x}(x, d) + d)$*.*

Proof. Consider any delay-sensitive utility function $U(x, d)$ that is strictly increasing in x for fixed d. Note that this implies $\forall x > 0, d > 0 : \frac{\partial U}{\partial x}(x, d) > 0$.

Suppose for some $c_1 > 0, \tau > 0$,

$$U(c_1, \tau) < \lim_{c_2 \to \infty} U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau) \tag{3.8}$$

Then in particular, for some $c_1 > 0, c_2 > 0, \tau > 0$,

$$U(c_1, \tau) < U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau) \tag{3.9}$$

Consider the following network.

Network 7:

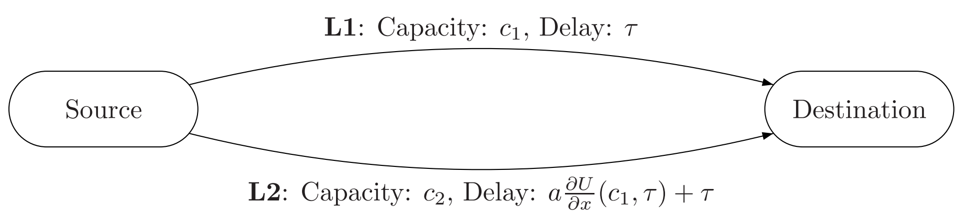

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x, d)$

Flow 1 achieves rate c_1 with propagation delay τ and utility $U(c_1, \tau)$. The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(c_1, \tau) + \tau$
R2	$a\frac{\partial U}{\partial x}(c_1, \tau) + \tau$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 1 had instead used route **R2**. Then it would achieve rate c_2 with propagation delay $a\frac{\partial U}{\partial x}(c_2, \tau) + \tau$ and utility $U(c_2, a\frac{\partial U}{\partial x}(c_2, \tau) + \tau)$.

From (3.9), we conclude that the equilibrium routing is suboptimal. □

Lemma 12 *Suppose that for all fixed $d > 0$, a delay-sensitive utility function $U(x, d)$ is strictly increasing in x. Further suppose that on all networks with sources using this utility function, if TCP/IP equilibrium exists, the equilibrium solves* (3.1) *and* (3.2)*. Then $\forall x > 0, d > 0 : U(x, d) > M(U, a\frac{\partial U}{\partial x}(x, d) + d)$.*

Proof. Consider any delay-sensitive utility function $U(x, d)$ that is strictly increasing in x for fixed d. Note that this implies $\forall x > 0, d > 0 : \frac{\partial U}{\partial x}(x, d) > 0$.

After using the previous lemma, we only have to show that $U(x, d) \neq M(U, a\frac{\partial U}{\partial x}(x, d) + d)$.

Suppose for some $c_1 > 0, \tau > 0$,

$$U(c_1, \tau) = \lim_{c_2 \to \infty} U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau) \tag{3.10}$$

Since $U(x, d)$ is strictly increasing for fixed d:

$$U(2c_1, \tau) > U(c_1, \tau)$$

Then for some $c_2 > 0$,

$$U(2c_1, \tau) - U(c_1, \tau) > U(c_1, \tau) - U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau) \tag{3.11}$$

Consider the following network.

Network 8:

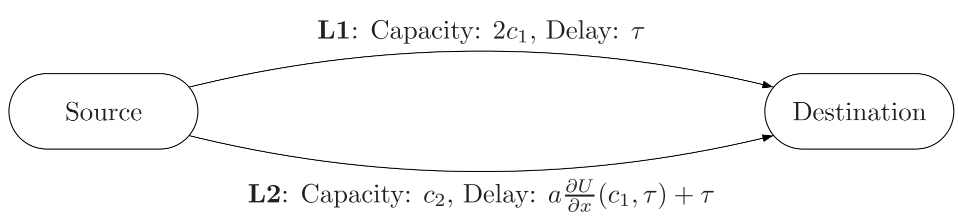

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**	**R1**	$U(x, d)$
Flow 2	**R1**, **R2**	**R1**	$U(x, d)$

Flow 1 and Flow 2 share link **L1** equally, both achieving rate c_1 with propagation delay τ and utility $U(c_1, \tau)$. Then the aggregate utility is $2U(c_1, \tau)$.

The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(c_1, \tau) + \tau$
R2	$a\frac{\partial U}{\partial x}(c_1, \tau) + \tau$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Suppose Flow 2 instead chose route **R2**. Flow 2 would achieve rate c_2 with propagation delay $a\frac{\partial U}{\partial x}(c_1, \tau) + \tau$ and utility $U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau)$. Flow 1 would achieve rate $2c_1$ with propagation delay τ and utility $U(2c_1, \tau)$. The aggregate utility achieved by this routing is:

$$U(2c_1, \tau) + U(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau)$$

From (3.11), we conclude that the equilibrium routing is suboptimal. □

Lemma 13 *Suppose that for all fixed $d > 0$, a delay-sensitive utility function $U(x, d)$ is strictly increasing in x, and $U(0, d)$ and $\frac{\partial U}{\partial x}(0, d)$ are both finite. Further suppose that on all networks with sources using this utility function, if TCP/IP equilibrium exists, the equilibrium solves* (3.1) *and* (3.2). *Then for all $\tau > 0$ so that $f(d) := M(U, d)$ is finite and continuous at $d = a\frac{\partial U}{\partial x}(0, \tau) + \tau$: $U(0, \tau) \leq M(U, a\frac{\partial U}{\partial x}(0, \tau) + \tau)$.*

Proof. Consider any delay-sensitive utility function $U(x, d)$ that is strictly increasing in x for fixed d. Note that this implies $\forall x > 0, d > 0 : \frac{\partial U}{\partial x}(x, d) > 0$.

Suppose there exists $\tau > 0$ so that $f(d) := M(U, d)$ is finite and continuous at $d = \frac{\partial U}{\partial x}(0, \tau) + \tau$, and:

$$U(0, \tau) > \lim_{c_2 \to \infty} U\left(c_2, a\frac{\partial U}{\partial x}(0, \tau) + \tau\right)$$

By continuity of $f(d)$ at $a\frac{\partial U}{\partial x}(0, \tau) + \tau$ and continuous differentiability of $U(x, d)$, $\exists c_1 > 0$ so that:

$$2U(\frac{3}{8}c_1, \tau) > \lim_{c_2 \to \infty} U\left(c_2, a\frac{\partial U}{\partial x}(c_1, \tau) + \tau\right) + U(\frac{3}{4}c_1, \tau)$$

For every $c_2 > 0$,

$$2U(\frac{3}{8}c_1,\tau) > U\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) + U(\frac{3}{4}c_1,\tau) \tag{3.12}$$

Also, we must have:

$$\lim_{c_2\to\infty} \frac{\partial U}{\partial x}\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) \geq 0$$

since the expression is monotonically decreasing in c_2 and bounded below by 0.

Suppose:

$$\lim_{c_2\to\infty} \frac{\partial U}{\partial x}\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) > 0$$

Then:

$$\lim_{c_2\to\infty} U\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) = \infty$$

But then $U(c_1,\tau)$ satisfies condition (3.8), and we can get a contradiction. So we assume:

$$\lim_{c_2\to\infty} \frac{\partial U}{\partial x}\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) = 0.$$

Since $U(x,d)$ is strictly concave increasing in x for fixed d,

$$\frac{\partial U}{\partial x}(c_1,\tau) < \frac{\partial U}{\partial x}(\frac{3}{4}c_1,\tau)$$

Combining these two facts, there exists $c_2 > 0$ so that:

$$a\frac{\partial U}{\partial x}(c_1,\tau) + a\frac{\partial U}{\partial x}\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right) + \tau < a\frac{\partial U}{\partial x}(\frac{3}{4}c_1,\tau)+\tau$$

Consider the following network.

Network 9:

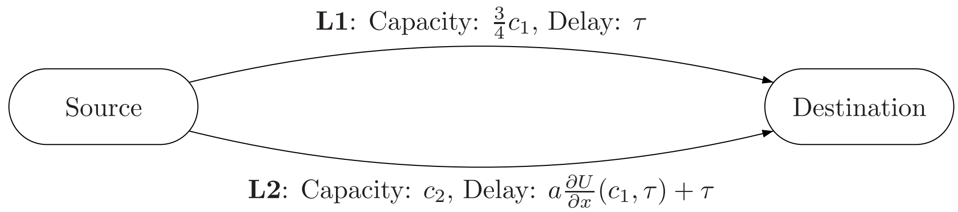

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**	**R1**	$U(x,d)$
Flow 2	**R1**, **R2**	**R2**	$U(x,d)$

Flow 1 achieves rate $\frac{3}{4}c_1$ with propagation delay τ and utility $U(\frac{3}{4}c_1,\tau)$. Flow 2 achieves a rate of c_2 with propagation delay $a\frac{\partial U}{\partial x}(c_1,\tau)+\tau$ and utility $U(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau)$.

The aggregate utility is:

$$U(c_2, \frac{\partial U}{\partial x}(c_1,\tau)+\tau) + U(\frac{3}{4}c_1,\tau)$$

The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(\frac{3}{4}c_1,\tau)+\tau$
R2	$a\frac{\partial U}{\partial x}(c_1,\tau)+a\frac{\partial U}{\partial x}\left(c_2, a\frac{\partial U}{\partial x}(c_1,\tau)+\tau\right)+\tau$

The initial routing is an equilibrium routing, since all flows are using minimal cost routes.

Consider the alternative routing where the second flow chooses route **R1**. Then both flows would share **L1** equally, and both would achieve rate $\frac{3}{8}c_1$ with utility $U(\frac{3}{8}c_1,\tau)$.

The aggregate utility of this alternative routing is:

$$2U(\frac{3}{8}c_1,\tau)$$

From (3.12), we conclude that the equilibrium routing is suboptimal. $\square$

Theorem 12 *Any class of delay-sensitive utility functions $\mathcal{B}$ such that when TCP/IP equilibrium exists, the equilibrium solves* (3.1) *and* (3.2) *for those utility functions, must have at least one of the following three properties:*

1. *$\exists U(x,d) \in \mathcal{B}, d > 0$ so that $U(x,d)$ is not strictly increasing in x.*

2. *$\forall U_1(x,d) \in \mathcal{B}$, $\forall \epsilon > 0$: $U_2(x,d) := U_1(x+\epsilon, d)$ is not in $\mathcal{B}$.*

3. *$\exists U(x,d) \in \mathcal{B}$, $D > 0$: $f(d) := M(U,d)$ is finite and discontinuous for all $d > D$.*

Proof. Suppose none of the properties hold. Then:

1. $\forall U(x,d) \in \mathcal{B}, d > 0$: $U(x,d)$ is strictly increasing in x.

2. $\exists U_1(x,d) \in \mathcal{B}$, $U_2(x,d) \in \mathcal{B}$, $\epsilon > 0$: $U_2(x,d) \equiv U_1(x+\epsilon, d)$.

3. $\forall U(x,d) \in \mathcal{B}$, $D > 0$: $f(d) := M(U,d)$ is either infinite or continuous for some $d > D$.

Note that for all $d > 0$, $U_2(0,d)$ and $\frac{\partial U_2}{\partial x}(0,d)$ are finite, since $U_1(\epsilon,d)$ and $\frac{\partial U_1}{\partial x}(\epsilon,d)$ must be finite.

Consider:

$$g(d) = a\frac{\partial U_2}{\partial x}(0,d) + d$$

Choose any $d_1 > 0$, and any $D > g(d_1)$. By assumption, there exists $d_2 > D$ so that $f(d) := M(U_2,d)$ is infinite or continuous at $d = d_2$. But by inspection, there exists d_3 so that $g(d_3) > d_2$. So $g(d_1) < d_2 < g(d_3)$. But $g(d)$ is continuous, so there exists d_4 so that $g(d_4) = d_2$. Lemma 12 implies that:

$$\begin{aligned} U_1(\epsilon, d_4) &> M(U_1, a\frac{\partial U_1}{\partial x}(\epsilon, d_4) + d_4) \\ U_2(0, d_4) &> M(U_2, a\frac{\partial U_2}{\partial x}(0, d_4) + d_4) \\ &> M(U_2, g(d_4)) \\ &> f(d_2) \end{aligned} \tag{3.13}$$

The first line follows from Lemma 12. The second line follows from the fact that $\frac{\partial U_2}{\partial x}(x,d) = \frac{\partial U_1}{\partial x}(x+\epsilon,d)$, and $M(U_1,d) = M(U_2,d)$.

So $f(d)$ is finite at d_2. By assumption, $f(d)$ must also be continuous at d_2, i.e. finite and continuous at $d_2 = g(d_4) = a\frac{\partial U_2}{\partial x}(0,d_4) + d_4$. But this, combined with (3.13), contradicts Lemma 13. $\square$

Remark 2 *We know that in some sense, we cannot improve on $\mathcal{C}$ by too much, and the proof techniques used to prove analogues of Theorem 1 and Theorem 2 do not seem to work for functions outside of $\mathcal{C}$. This suggests that $\mathcal{C}$ is likely to be the only full class of functions that TCP/IP optimizes.*

Chapter 4

Stability and Utility of Routing Policies

In this chapter, we analyze the effects on stability and utility that result from adjusting routing policy.

4.1 Stability

4.1.1 Definitions

Denote the set of paths $\mathcal{G}^i \subseteq \mathcal{H}^i$ with minimal propagation delay for flow i by:

$$\mathcal{G}^i := \left\{ r^i \in \mathcal{H}^i : \tau^T r^i = \min_{s^i \in \mathcal{H}^i} \tau^T s^i \right\}$$

Denote the set of paths $\mathcal{F}^i \subseteq \mathcal{H}^i$ without minimal propagation delay for flow i by:

$$\mathcal{F}^i := \mathcal{H}^i - \mathcal{G}^i$$

Define $q(R)$, a function that computes the equilibrium congestion price vector for a given routing matrix $R \in \mathcal{R}_s$. We assume that it implicitly depends on an arbitrary, fixed network $(L, N, \mathcal{F}^i, \mathcal{G}^i, \mathcal{H}^i, \mathcal{R}_s, K^i, U_i, \tau, c)$:

$$q(R) = \arg \min_{p \geq 0} \max_{x \geq 0} \left(\sum_i^N U_i(x_i, d_i) - p^T R x \right) + p^T c$$

4.1.2 Mostly-static-cost routing on networks

4.1.2.1 Effect of small a on path choices

In this section, we show that with sufficiently small a, all flows will choose only minimal propagation delay paths.

Define $h(x, y)$, a function that will be used to simplify notation:

$$h(x, y) = \begin{cases} \frac{x}{y} & \text{if } y > 0 \\ \infty & \text{if } y \leq 0 \end{cases}$$

For notational simplicity, all of the following functions, lemmas, and theorems in this section implicitly depend on an arbitrary, fixed network $(L, N, \mathcal{F}^i, \mathcal{G}^i, \mathcal{H}^i, \mathcal{R}_s, K^i, U_i, \tau, c)$.

Define $a_{\#}$ as follows:

$$a_{\#} = \min_{R \in \mathcal{R}_s} \min_{0 < i \leq N} \begin{cases} \max_{m^i \in \mathcal{G}^i} \min_{r^i \in \mathcal{F}^i} h\left(\tau^T(r^i - m^i), q(R)^T(m^i - r^i)\right) & \text{if } \mathcal{F}^i \neq \emptyset \\ \infty & \text{if } \mathcal{F}^i = \emptyset \end{cases}$$

Lemma 14 *$a_{\#}$ is strictly positive.*

Proof. This is clear by inspection. □

Theorem 13 *Suppose $a < a_{\#}$. Then:*

$$\forall R \in \mathcal{R}_s, \forall 0 < i \leq N, \exists m^i \in \mathcal{G}^i, \forall r^i \in \mathcal{F}^i : (aq(R) + \tau)^T m^i < (aq(R) + \tau)^T r^i. \tag{4.1}$$

In other words, all flows on the network will choose a path with minimal propagation delay in the next iteration.

Proof. We manipulate the right hand side of (4.1).

$$\begin{aligned} (aq(R) + \tau)^T m^i &< (aq(R) + \tau)^T r^i \\ aq(R)^T m^i + \tau^T m^i &< aq(R)^T r^i + \tau^T r^i \\ aq(R)^T m^i - aq(R)^T r^i &< \tau^T r^i - \tau^T m^i \\ aq(R)^T (m^i - r^i) &< \tau^T (r^i - m^i) \end{aligned}$$

But by inspecting the definition of $h(x, y)$, this inequality holds if

$$a < h\left(\tau^T(r^i - m^i), q(R)^T(m^i - r^i)\right),$$

since $\tau^T(r^i - m^i) > 0$ if $m^i \in \mathcal{G}^i$ and $r^i \in \mathcal{F}^i$.

The formal part of the theorem is then easy to see. It implies that for any current routing, and for every flow, a path with minimal propagation delay has strictly lower cost than any path without minimal propagation delay. Therefore no flow will select any path without minimal propagation delay. Therefore every flow will select a path with minimal propagation delay. □

4.1.2.2 Networks with unique minimal propagation delay paths

Theorem 14 *Suppose all source-destination pairs on a network have unique minimum-propagation-delay paths. Then if $a < a_{\#}$, TCP/IP has an asymptotically stable equilibrium.*

Proof. Each flow only has one path with minimal propagation delay. Applying Theorem 13, each flow will always select the same path, and it will always do this from any routing. □

Corollary 1 *Suppose every path in a network has different propagation delay. Then if $a < a_{\#}$, TCP/IP has an asymptotically stable equilibrium.*

Proof. If every path in the network has different propagation delay, every source-destination pair on a network has a unique minimum-propagation-delay path, so the result of Theorem 14 applies. □

4.1.2.3 Networks without unique minimal propagation delay paths

Not all networks can be stabilized by decreasing a. Consider the following network, assuming that all sources use an arbitrary delay-insensitive utility function $U(x)$.

Network 10:

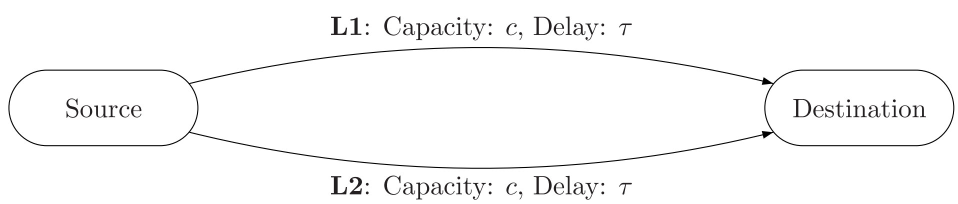

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**	**R1**	$U(x)$

Remark 3 *For all $a > 0$, the above network has no equilibrium.*

Proof. Flow 1 achieves rate c, with utility $U(c)$. The route costs are then:

Route Costs	
Route	Cost
R1	$aU'(c) + \tau$
R2	τ

For every $a > 0$, Flow 1 will choose **R2** in the next iteration.

But then Flow 1 achieves rate c with utility $U(c)$. The route costs are then:

Route Costs	
Route	Cost
R1	τ
R2	$aU'(c) + \tau$

For every $a > 0$, Flow 1 will choose **R1** in the next iteration.

So the routing will oscillate between **R1** and **R2**. It is also easy to see that this would happen even if the flow was initially on route **R2**. Therefore this network has no equilibrium. $\square$

We now generalize this example. We know that $\exists a_{\#}$ so that if $a < a_{\#}$, all flows choose between paths with minimal propagation delay. When there are multiple such paths available, there may or may not be an equilibrium.

Lemma 15 *Suppose we have some network, and routing policy on this network is such that $a < a_{\#}$. Then for any attainable price vectors p so that $p = q(R)$ for some $R \in \mathcal{R}_s$, and for all $0 < i \leq N$,*

$$p^T r^i = \min_{s^i \in \mathcal{G}^i} p^T s^i, \ \textit{where } r^i \in \mathcal{G}^i \tag{4.2}$$

if and only if

$$(ap + \tau)^T r^i = \min_{s^i \in \mathcal{H}^i} (ap + \tau)^T s^i, \ \textit{where } r^i \in \mathcal{H}^i \tag{4.3}$$

Proof. First, suppose (4.3) holds.

From Theorem 14, we know that:

$$\exists m^i \in \mathcal{G}^i, \forall r^i \in (\mathcal{H}^i - \mathcal{G}^i) : (ap + \tau)^T m^i < (ap + \tau)^T r^i.$$

So:

$$(ap+\tau)^T r^i = \min_{s^i \in \mathcal{G}^i} (ap+\tau)^T s^i$$

and r^i must be in $\mathcal{G}^i$. But since every path in $\mathcal{G}^i$ has minimal propagation delay:

$$\forall r^i \in \mathcal{G}^i, s^i \in \mathcal{G}^i : \tau^T r^i = \tau^T s^i$$

So:

$$(ap)^T r^i = \min_{s^i \in \mathcal{G}^i} (ap)^T s^i,$$

But then:

$$p^T r^i = \min_{s^i \in \mathcal{G}^i} p^T s^i,$$

as desired. For the reverse case, we use the above argument backwards, and note that $\mathcal{G}^i \subseteq \mathcal{H}^i$.

□

Theorem 15 *Suppose we have some network, and routing policy on this network is such that $a < a_{\#}$. Consider the modified network obtained by deleting all paths without minimal propagation delay from the original network. Then the original network with routing based on $ap_l + d_l$ has the same equilibrium and stability properties as the modified network with routing based on p_l.*

Proof. Consider the TCP/IP dynamical system on the modified network when link costs are p_l:

$$(p(t))^T r^i(t+1) = \min_{r^i \in \mathcal{G}^i} (p(t))^T r^i, \text{ for all } i. \ (\forall i, t\colon r^i(t) \in \mathcal{G}^i)$$

$$(p(t), x(t)) = \arg \min_{p \geq 0} \max_{x \geq 0} \left(\sum_i^N U_i(x_i, d_i) - p^T R(t) x \right) + p^T c$$

Consider the TCP/IP dynamical system on the original network when link costs are $ap_l + \tau_l$:

$$(ap(t)+\tau)^T r^i(t+1) = \min_{r^i \in \mathcal{H}^i} (ap(t)+\tau)^T r^i, \text{ for all } i. \ (\forall i, t\colon r^i(t) \in \mathcal{H}^i)$$

$$(p(t), x(t)) = \arg \min_{p \geq 0} \max_{x \geq 0} \left(\sum_i^N U_i(x_i, d_i) - p^T R(t) x \right) + p^T c$$

But the previous lemma implies that these dynamical systems are equivalent. Since they are equivalent, they share the same equilibrium and stability properties. □

4.1.3 Destabilization from increased static component in link cost

In this section, we will show that it is possible to destabilize a network by increasing the static component in link cost. This seems counterintuitive, but is true.

Suppose $U(x, d)$ is an arbitrary delay-sensitive or delay-insensitive function.

Consider the following inequalities:

$$\frac{\partial U}{\partial x}(c_1, \tau_1) > \frac{\partial U}{\partial x}(c_2, \tau_2) > 0 \tag{4.4}$$

$$c_2 > c_1 > 0 \tag{4.5}$$

$$\tau_2 > \tau_1 > 0 \tag{4.6}$$

Lemma 16 *If $U(x, d)$ is delay-insensitive, it is possible to choose c_1, c_2, τ_1, τ_2 so that they satisfy equations* (4.4) *through* (4.6)*.*

Proof. Suppose we have a delay-insensitive utility function $U(x, d) \equiv U(x)$. Choose any $c_2 > c_1 > 0$, and any $\tau_2 > \tau_1 > 0$. Then we need to show that $U'(c_1) > U'(c_2) > 0$. But this is true since $U(x)$ is strictly concave increasing. □

Lemma 17 *If $U(x, d)$ is delay-sensitive, it is possible to choose c_1, c_2, τ_1, τ_2 so that they satisfy equations* (4.4) *through* (4.6)*.*

Proof.

Suppose we have a delay-sensitive utility function $U(x, d)$. From Theorem 5, there exist $\tau_1 > 0, c_2 > c_1 > 0$ so that:

$$\frac{\partial U}{\partial x}(c_1, \tau_1) > \frac{\partial U}{\partial x}(c_2, \tau_1) > 0.$$

But since $U(x, d)$ is continuously differentiable, there exists $\epsilon > 0$ so that:

$$\frac{\partial U}{\partial x}(c_1, \tau_1) > \frac{\partial U}{\partial x}(c_2, \tau_1 + \epsilon) > 0.$$

Choose $\tau_2 = \tau_1 + \epsilon$. Evidently, equations (4.4) through (4.6) hold with this choice of c_1, τ_1, c_2, τ_2 as desired. □

Suppose c_1, τ_1, c_2, τ_2 are chosen so that they satisfy equations (4.4) through (4.6). Consider the following network.

Network 11:

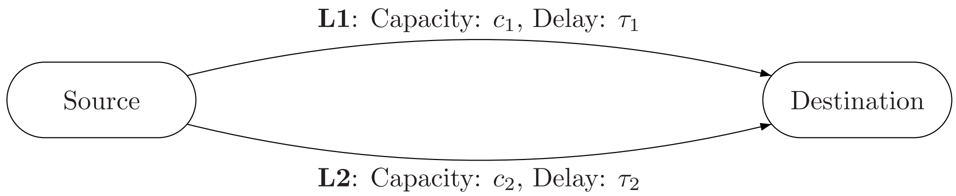

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**	**R1**	$U(x,d)$
Flow 2	**R1**, **R2**		$U(x,d)$

Define function $A_3(U, c_1, \tau_1, c_2, \tau_2)$, where U is a delay-insensitive or delay-sensitive utility function and the rest of the parameters are in $\mathbb{R}$ as follows:

$$A_3(U, c_1, \tau_1, c_2, \tau_2) = \frac{\tau_2 - \tau_1}{\frac{\partial U}{\partial x}(c_1, \tau_1) - \frac{\partial U}{\partial x}(c_2, \tau_2)}$$

Lemma 18 *Consider any delay-sensitive or delay-insensitive utility function $U(x,d)$. Suppose c_1, τ_1, c_2, τ_2 satisfy (4.4) through (4.6) with $U(x,d)$. Network 11 with all sources using $U(x,d)$ is stable for all $a > A_3(U, c_1, \tau_1, c_2, \tau_2)$.*

Proof. We show that routing converges from every possible initial condition.

Suppose Flow 2 is on route **R1**. Then Flow 1 and Flow 2 share **L1** equally, and both achieve rate $\frac{c_1}{2}$ with propagation delay τ_1 and utility $U(\frac{c_1}{2}, \tau_1)$.

The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1) + \tau_1$
R2	τ_2

We verify that if $a > A_3(U, c_1, \tau_1, c_2, \tau_2)$, then **R2** is a lower cost route than **R1**:

$$
\begin{aligned}
a\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)+\tau_1 &> \frac{\tau_2-\tau_1}{\frac{\partial U}{\partial x}(c_1,\tau_1)-\frac{\partial U}{\partial x}(c_2,\tau_2)}\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)+\tau_1\\
&> \frac{\tau_2-\tau_1}{\frac{\partial U}{\partial x}(c_1,\tau_1)}\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)+\tau_1\\
&> \frac{\tau_2-\tau_1}{\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)}\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)+\tau_1\\
&> \tau_2-\tau_1+\tau_1\\
&> \tau_2
\end{aligned}
$$

To show asymptotic stability, it is then sufficient to show that the routing where Flow 2 is on **R2** is an equilibrium.

Suppose Flow 2 is on **R2**. Then Flow 1 achieves rate c_1 with propagation delay τ_1 and utility $U(c_1,\tau_1)$, and Flow 2 achieves rate c_2 with propagation delay τ_2 and utility $U(c_2,\tau_2)$. The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(c_1,\tau_1)+\tau_1$
R2	$a\frac{\partial U}{\partial x}(c_2,\tau_2)+\tau_2$

Consider Flow 2's route choice at the next routing iteration.

We verify that if $a > A_3(U,c_1,\tau_1,c_2,\tau_2)$, then **R2** is a lower cost route than **R1**:

$$
\begin{aligned}
a &> \frac{\tau_2-\tau_1}{\frac{\partial U}{\partial x}(c_1,\tau_1)-\frac{\partial U}{\partial x}(c_2,\tau_2)}\\
a(\frac{\partial U}{\partial x}(c_1,\tau_1)-\frac{\partial U}{\partial x}(c_2,\tau_2)) &> \tau_2-\tau_1\\
a\frac{\partial U}{\partial x}(c_1,\tau_1)+\tau_1 &> a\frac{\partial U}{\partial x}(c_2,\tau_2)+\tau_2
\end{aligned}
$$

Therefore **R2** is minimal cost, and so this routing is an equilibrium routing. □

Define function $A_2(U,c_1,\tau_1,c_2,\tau_2)$, where U is a delay-insensitive or delay-sensitive utility function and the rest of the parameters are in $\mathbb{R}$ as follows:

$$
A_2(U,c_1,\tau_1,c_2,\tau_2)=\frac{\tau_2-\tau_1}{\frac{\partial U}{\partial x}(\frac{c_1}{2},\tau_1)}
$$

Lemma 19 *Consider any delay-sensitive or delay-insensitive utility function* $U(x,d)$. *Suppose* c_1, τ_1, c_2, τ_2 *satisfy* (4.4) *through* (4.6) *with* $U(x,d)$. *Then* $A_2(U, c_1, \tau_1, c_2, \tau_2) < A_3(U, c_1, \tau_1, c_2, \tau_2)$.

Proof.

$$\begin{aligned}
A_3(\ldots) - A_2(\ldots) &= \frac{\tau_2 - \tau_1}{\frac{\partial U}{\partial x}(c_1, \tau_1) - \frac{\partial U}{\partial x}(c_2, \tau_2)} - \frac{\tau_2 - \tau_1}{\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1)} \\
&= (\tau_2 - \tau_1)\left(\frac{1}{\frac{\partial U}{\partial x}(c_1, \tau_1) - \frac{\partial U}{\partial x}(c_2, \tau_2)} - \frac{1}{\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1)}\right) \\
&> (\tau_2 - \tau_1)\left(\frac{1}{\frac{\partial U}{\partial x}(c_1, \tau_1)} - \frac{1}{\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1)}\right) \\
&> (\tau_2 - \tau_1)\left(\frac{\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1) - \frac{\partial U}{\partial x}(c_1, \tau_1)}{\frac{\partial U}{\partial x}(c_1, \tau_1)\frac{\partial U}{\partial x}(c_2, \tau_2)}\right) \\
&> 0
\end{aligned}$$

□

Lemma 20 *Consider any delay-sensitive or delay-insensitive utility function* $U(x,d)$. *Suppose* c_1, τ_1, c_2, τ_2 *satisfy* (4.4) *through* (4.6) *with* $U(x,d)$. *Network 11 with all sources using* $U(x,d)$ *is unstable for all* a *satisfying* $A_2(U, c_1, \tau_1, c_2, \tau_2) < a < A_3(U, c_1, \tau_1, c_2, \tau_2)$.

Proof. Suppose Flow 2 is on **R2**. Then Flow 1 achieves rate c_1 with propagation delay τ_1 and utility $U(c_1, \tau_1)$, and Flow 2 achieves rate c_2 with propagation delay τ_2 and utility $U(c_2, \tau_2)$. The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(c_1, \tau_1) + \tau_1$
R2	$a\frac{\partial U}{\partial x}(c_2, \tau_2) + \tau_2$

Consider Flow 2's route decision at the next routing iteration.

$$\begin{aligned}
a &< A_3(U, c_1, \tau_1, c_2, \tau_2) \\
a &< \frac{\tau_2 - \tau_1}{\frac{\partial U}{\partial x}(c_1, \tau_1) - \frac{\partial U}{\partial x}(c_2, \tau_2)} \\
a(\frac{\partial U}{\partial x}(c_1, \tau_1) - \frac{\partial U}{\partial x}(c_2, \tau_2)) &< \tau_2 - \tau_1 \\
a\frac{\partial U}{\partial x}(c_1, \tau_1) + \tau_1 &< a\frac{\partial U}{\partial x}(c_2, \tau_2) + \tau_2
\end{aligned}$$

So Flow 2 next chooses route **R1**. Then Flow 1 and Flow 2 share **L1** equally, and both achieve rate $\frac{c_1}{2}$ with delay τ_1 and utility $U(\frac{c_1}{2}, \tau_1)$. The route costs are then:

Route Costs	
Route	Cost
R1	$a\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1) + \tau_1$
R2	τ_2

But:

$$\begin{aligned} a &> A_2(U, c_1, \tau_1, c_2, \tau_2) \\ a\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1) + \tau_1 &> \frac{\tau_2 - \tau_1}{\frac{\partial U}{\partial x}(\frac{c_1}{2}, \tau_1)} \frac{\partial U}{\partial x}(\frac{c_1}{2}, c_1) + \tau_1 \\ &> \tau_2 \end{aligned}$$

So Flow 2 next chooses route **R2**.

This implies that Flow 2's routing oscillates between **R1** and **R2**, so the network is unstable. □

Theorem 16 *Consider any delay-sensitive or delay-insensitive utility function $U(x, d)$. There exist a network with sources using this utility function, and constants $a_3 > a_2 \geq a_1 > 0$ so that:*

1. *The network is stable for $a \in (0, a_1)$.*
2. *The network is unstable for $a \in (a_2, a_3)$.*
3. *The network is stable for $a \in (a_3, \infty)$.*

Proof. We choose c_1, τ_1, c_2, τ_2 so that they satisfy (4.4) through (4.6). Then consider Network 11 with all sources using $U(x, d)$. Set $a_1 = a_\#$ for this network. Theorem 13 implies that for $a < a_1$, the network is stable. We then set $a_2 = A_2(U, c_1, \tau_1, c_2, \tau_2)$ and $a_3 = A_3(U, c_1, \tau_1, c_2, \tau_2)$. The lemmas in this section then establish the desired result. □

4.2 Utility

[5] analyzed the effects of increasing a on time-averaged aggregate utility for a ring network, and a randomly generated network.

On the ring network, time-averaged aggregate utility approached the maximum possible time-averaged aggregate utility for any a, as a increased. On the generated network, time-averaged aggregate utility increased until routing stability set in, and then decreased.

In this section, we show that the effects of increasing a on time-averaged aggregate utility are extremely network dependent.

In particular, we show how to construct a network with any given utility profile as a function of the weight a.

4.2.1 Effect of increasing a on utility

For every delay-insensitive $U(x)$, there exist $k^+ > k^* > k^- > 0$ and $g > 0$ so that $g = U(k^+) - U(k^*) = U(k^*) - U(k^-)$. Also, for every $z \in [-1, 1]$, there exists $k \in [k^-, k^+]$ so that $U(k) - U(k^*) = zg$, that is given by $k(z) := U^{-1}(U(k^*) + zg)$. This is easy to see since $U(x)$ is strictly monotone increasing.

Consider the following network $N(j, z)$, parameterized by $j > 0$, and $z \in [-1, 1]$:

Network 12:

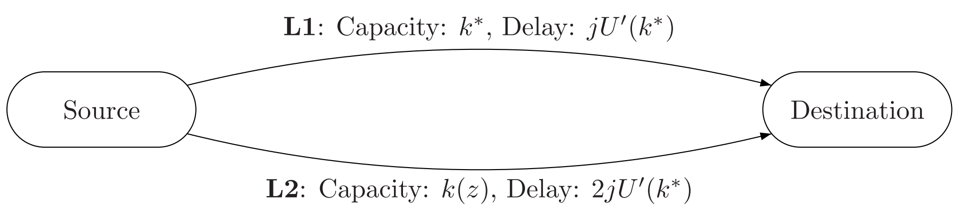

Network Routes	
Route	Path
R1	**L1**
R2	**L2**

Flows			
Flow	Possible Routes	Initial Route	Utility Function
Flow 1	**R1**, **R2**		$U(x)$

Denote the time-averaged utility of the network under routing policy $ap_l + d_l$ by $T(N, a)$.

Lemma 21 *For every $j > 0$, $z \in [-1, 1]$, network $N(j, z)$ has the following properties:*

1. $\forall a_1 \in [0, j], a_2 \in [0, j]$: $T(N, a_1) = T(N, a_2)$

2. $\forall a_1 \in (j, \infty), a_2 \in (j, \infty)$: $T(N, a_1) = T(N, a_2)$

3. $\forall a_1 \in (0, j], a_2 \in (j, \infty)$: $T(N, a_2) = T(N, a_1) + \frac{zg}{2}$

Proof. Suppose current routing is **R2**. Then the route costs are:

Route Costs	
Route	Cost
R1	$jU'(k^*)$
R2	$2jU'(k^*) + aU'(k(z))$

By inspection, for every $a \geq 0$, **R1** has less cost than **R2**. Therefore routing on the next iteration will be **R1**.

Suppose current routing is **R1**. Then the route costs are:

Route Costs	
Route	Cost
R1	$jU'(k^*) + aU'(k^*)$
R2	$2jU'(k^*)$

By inspection, if $a \leq j$ then **R1** is a lower cost route than **R2**, and if $a > j$, then **R2** is a strictly lower cost route than **R1**.

Flow 1's utility on route **R1** between routing changes is $U(k^*)$. Flow 1's utility on route **R2** between routing changes is $U(kz)$.

If $a \leq j$, then the network is stable with Flow 1 on **R1** and the aggregate utility is $U(k^*)$. If $a > j$, then the network oscillates between **R1** and **R2** and achieves time-averaged utility $\frac{U(k(z))+U(k^*)}{2}$. The excess utility gained by increasing a from less than j to more than j is given by:

$$\frac{U(k(z)) + U(k^*)}{2} - U(k^*) = \frac{U(k(z)) - U(k^*)}{2} = \frac{zg}{2}$$

□

Definition 3 *A utility-versus-a profile is a pair of vectors (x, y) such that $|x| = |y| > 0$, $\forall i$: $x_i > 0$, and $\forall i < |x|$: $x_i < x_{i+1}$.*

Definition 4 *A network N matches a utility-versus-a profile (x, y) if there exists $c > 0$ so that:*

If $|x| > 1$,

1. $\forall 1 < i < |x|$, $\forall a_1$ *s.t.* $x_{i-1} < a_1 \leq x_i$, $\forall a_2$ *s.t.* $x_i < a_2 < x_{i+1}$: $T(N, a_2) - T(N, a_1) = cy_i$.
2. $\forall a_1$ *s.t.* $0 < a_1 \leq x_1$, $\forall a_2$ *s.t.* $x_1 < a_2 < x_2$: $T(N, a_2) - T(N, a_1) = cy_1$.
3. $\forall a_1$ *s.t.* $x_{|x|-1} < a_1 \leq x_{|x|}$, $\forall a_2$ *s.t.* $x_{|x|} < a_2 < \infty$: $T(N, a_2) - T(N, a_1) = cy_{|x|}$.

If $|x| = 1$, $\forall a_1$ *s.t.* $0 < a_1 \leq x_1$, $\forall a_2$ *s.t.* $x_1 < a_2 < \infty$*:* $T(N, a_2) - T(N, a_1) = cy_1$*.*

In other words, for all i*, the time-averaged aggregate utility of the network increases by* cy_i *at* $a = x_i$*.*

Theorem 17 *For every utility-versus-a profile* (x, y)*, there exists a network with sources using delay-insensitive utility functions that matches this profile.*

Proof. Consider any utility-versus-a profile (x, y). Define the normalized y as $y^* := |\max_i y_i|^{-1} y$. Construct the network N^* given by taking the union of networks $N_1 \ldots N_{|y|}$ where $\forall i$, $N_i := N(x_i, y_i^*)$. (The subnetworks are entirely disjoint in the union network).

It is easy to see that the network N^* matches profile (x, y) with $c = \frac{g}{2|\max_i v_i|}$. $\square$

Chapter 5

Conclusion and Future Work

In this thesis, we analyzed whether or not TCP/IP implicitly maximizes delay-sensitive or delay-insensitive utility functions with different link cost definitions. We identified a class of delay-sensitive utility functions $\mathcal{C}$ that TCP/IP does implicitly maximize with link cost $ap_l + d_l$, and proved some general results about any class that has the same property.

We further analyzed path choice and stability properties when the weight on the congestion price is sufficiently small. We also showed that one can construct a network with any given utility profile as a function of the weight a.

Our results have not completely established that there are no other classes of utility functions $\mathcal{B}$ that TCP/IP implicitly maximizes with link cost $ap_l + d_l$. Further, our results do not fully explain the stability and utility properties of general a on general networks. We plan further study in these directions.

Bibliography

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

[2] 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, 1998.

[3] Steven H. Low. A duality model of tcp and queue management algorithms. *IEEE/ACM Trans. Netw.*, 11(4):525–536, 2003.

[4] Steven H. Low, Larry L. Peterson, and Limin Wang. Understanding tcp vegas: a duality model. *J. ACM*, 49(2):207–235, 2002.

[5] Jiantao Wang, Lun Li, Steven H. Low, and John C. Doyle. Cross-layer optimization in tcp/ip networks. *IEEE/ACM Trans. Netw.*, 13(3):582–595, 2005.