

Heterogeneous Congestion Control Protocols

Thesis by
Ao (Kevin) Tang

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2006
(Defended May 3, 2006)

Acknowledgements

I want to deeply thank my advisor, Professor Steven H. Low, for his guidance and support. It has been very amazing for me since the beginning to see how Steven combines keen enthusiasm on very theoretical technical details on broad topics with serious interest in a wide range of experiments and real applications. That requires not only a lot of vision and energy, but also courage. As an advisor, Steven has always maintained high standards for us students and has been doing a lot to help us grow. I feel really lucky in working with Steven during my formative years and I don't think I can ask for more from an advisor.

My gratitude also extends to Professor Jehoshua Bruck, Mani Chandy, John Doyle and Babak Hassibi for serving in my committee for both the candidacy and the final exam. I benefit a lot from our discussions on research and career development.

Special thanks go to Professor Mung Chiang at Princeton and Professor Asuman Ozdaglar and Daron Acemoglu at MIT. Thank Mung for our fruitful collaborations and letting me know what hard working is. Thank Asu and Daron for working with me on fun problems, showing me the beauty of theory and the devotion such beautiful things deserve to have.

I have benefited enormously from my colleagues and friends at Caltech and also at Princeton and MIT. I especially want to thank Dr. Jiantao Wang, David Wei and Alp Simsek whose collaborations and discussions with me directly helped development of this thesis and Dr. Cheng Jin for being such a good and supportive friend for years. Finally, I have been extremely fortunate to have Christine Ortega as our secretary. She helps arrange things into every detail and always perfectly.

Last but not least, I want to thank my parents for their unconditional support and love and thank Lu Tian for the happy days. These days coincide with the productive period when this thesis is taking shape and I don't treat it as a pure coincidence.

Abstract

Homogeneity of price is an implicit yet fundamental assumption underlying price based resource allocation theory. In this thesis, we study the effects of relaxing this assumption by examining a concrete engineering system (network with heterogeneous congestion control protocols). The behavior of the system turns out to be very different from the homogeneous case and can potentially be much more complicated. A systematic theory is developed that includes all major properties of equilibrium of the system such as existence, uniqueness, optimality, and stability. In addition to analysis, we also present numerical examples, simulations, and experiments to illustrate the theory and verify its predictions.

When heterogeneous congestion control protocols that react to different pricing signals share the same network, the resulting equilibrium can no longer be interpreted as a solution to the standard utility maximization problem as the current theory suggests. After introducing a mathematical formulation of network equilibrium for multi-protocol networks, we prove the existence of equilibrium under mild assumptions. For almost all networks, the equilibria are locally unique. They are finite and odd in number. They cannot all be locally stable unless the equilibrium is globally unique. We also derive two conditions for global uniqueness. By identifying an optimization problem associated with *every* equilibrium, we show that every equilibrium is Pareto efficient and provide an upper bound on efficiency loss due to pricing heterogeneity. Both intra-protocol and inter-protocol fairness are then discussed. On dynamics, various stability results are provided. In particular it is shown that if the degree of pricing heterogeneity is small enough, the network equilibrium is not only unique but also locally stable. Finally, a distributed algorithm is proposed to steer a network to the unique equilibrium that maximizes the aggregate utility, by only updating a linear parameter in the sources' algorithms in a slow timescale.

Contents

| | |
|---|------------|
| Acknowledgements | iii |
| Abstract | iv |
| 1 Introduction | 1 |
| 1.1 Network Congestion Control | 2 |
| 1.2 A Motivating Example | 3 |
| 1.3 Contributions of This Thesis | 9 |
| 2 Existence and Uniqueness | 11 |
| 2.1 Model and Notations | 12 |
| 2.1.1 Notations | 12 |
| 2.1.2 Network equilibrium | 13 |
| 2.1.3 Current theory: $J = 1$ | 15 |
| 2.2 Existence | 17 |
| 2.3 Uniqueness | 19 |
| 2.3.1 Multiple equilibria: examples | 19 |
| 2.3.2 Regular networks | 25 |
| 2.3.3 Proofs | 30 |
| 2.3.4 Global uniqueness: mapping functions $m(p)$ | 33 |
| 2.3.4.1 General result | 34 |
| 2.3.4.2 $L = 3$ and $J = 2$ case | 35 |
| 2.3.5 Global uniqueness: Jacobian $\mathbf{J}(p)$ | 37 |
| 2.3.6 Global uniqueness of special networks | 41 |

| | | |
|----------|--|-----------|
| 2.3.6.1 | Case 1: linear link-independent m^j | 41 |
| 2.3.6.2 | Case 2: linear network | 42 |
| 2.3.6.3 | Case 3: networks with no flow using more than two links | 45 |
| 2.4 | Related Work in Economics | 47 |
| 3 | Optimality and Stability | 50 |
| 3.1 | Optimality | 51 |
| 3.1.1 | Efficiency | 51 |
| 3.1.1.1 | Qualitative results: Pareto efficiency | 51 |
| 3.1.1.2 | Quantitative results: price of heterogeneity | 52 |
| 3.1.2 | Fairness | 54 |
| 3.1.2.1 | Intra-protocol fairness | 54 |
| 3.1.2.2 | Inter-protocol fairness | 55 |
| 3.2 | Stability | 58 |
| 3.2.1 | Low dimensional results | 58 |
| 3.2.2 | General case | 60 |
| 4 | A Slow Timescale Update | 62 |
| 4.1 | Two Examples | 63 |
| 4.1.1 | Dependence of bandwidth allocation on buffer size | 63 |
| 4.1.2 | Dependence of bandwidth allocation on flow arrival pattern | 64 |
| 4.2 | One Solution | 67 |
| 4.2.1 | Independence of bandwidth allocation on buffer size | 68 |
| 4.2.2 | Independence of bandwidth allocation on flow arrival pattern | 69 |
| 4.3 | Analysis | 72 |
| 4.4 | Numerical Examples | 75 |
| 4.5 | WAN in Lab Experiments | 79 |
| 5 | Conclusion | 82 |

| | |
|---|----------------|
| 6 Appendix | 87 |
| 6.1 Simulation of Multiple Equilibria | 87 |
| 6.2 Proof of Theorem 2.8 | 90 |
| 6.3 Proof of Theorem 2.10 | 97 |
| 6.4 Smallest Network with Multiple Equilibria | 100 |
| 6.5 Proof of Lemma 2.16 | 101 |
| 6.6 Proof of Theorem 3.8 | 102 |
| 6.7 Proof of Theorem 4.2 | 108 |
| Bibliography | 110 |

List of Figures

| | | |
|------|---|----|
| 1.1 | Dummysnet setup for Experiments 1.1 and 1.2. | 5 |
| 1.2 | Multiple equilibria scenario. | 6 |
| 1.3 | Experiment 1.1: queue size and aggregate throughput at link 1. | 7 |
| 1.4 | Experiment 1.1: queue size and aggregate throughput at link 2. | 8 |
| 2.1 | Example 2.1: uncountably many equilibria. | 21 |
| 2.2 | Two equilibrium with different active bottleneck link sets. | 22 |
| 2.3 | Example 2.3: construction of multiple isolated equilibria. | 29 |
| 2.4 | Example 2.3: vector field of (p_1, p_2) | 30 |
| 2.5 | Region of λ_l for global uniqueness: projection to $\lambda_1 - \lambda_2$ plane. | 36 |
| 2.6 | Region of λ_l for global uniqueness: cross-section cut by plane $\lambda_1 + \lambda_2 = 1$ | 37 |
| 2.7 | Theorem 2.15: linear network. | 43 |
| 2.8 | Experiment 2.1: unique equilibrium. | 45 |
| 2.9 | Experiment 2.1: queue size and aggregate throughput at link 1. | 45 |
| 2.10 | Experiment 2.1: queue size and aggregate throughput at link 2. | 46 |
| 2.11 | Experiment 2.1: queue size and aggregate throughput at link 3. | 46 |
| 4.1 | Single link example. | 64 |
| 4.2 | FAST vs. Reno with a buffer size of 400 pkts. | 65 |
| 4.3 | FAST vs. Reno with a buffer size of 80 pkts. | 65 |
| 4.4 | Multiple equilibria scenario. | 66 |
| 4.5 | Bandwidth shares of Reno and FAST when FAST starts first. | 67 |
| 4.6 | Bandwidth shares of Reno and FAST when Reno starts first. | 67 |
| 4.7 | FAST vs. Reno, with buffer size of 400 pkts and Algorithm 1 | 69 |

| | | |
|------|---|----|
| 4.8 | FAST vs. Reno, with buffer size of 80 pkts and Algorithm 1 | 70 |
| 4.9 | α trajectory in experiment 4.1b | 70 |
| 4.10 | FAST starts first with Algorithm 1. | 71 |
| 4.11 | Reno starts first with Algorithm 1. | 71 |
| 4.12 | α trajectory in experiment 4.2b | 72 |
| 4.13 | Case 1: $p_1(t)$ with different κ_i^j | 78 |
| 4.14 | Case 2: $p_1(t)$ with different κ_i^j | 78 |
| 4.15 | $p(t)$ with different probability of updating | 79 |
| 4.16 | WAN in Lab experiment setup | 80 |
| 4.17 | Bandwidth partition between Reno and FAST | 81 |
| 4.18 | Bandwidth sharing among FAST flows | 81 |
| 6.1 | Experiment 6.1: Aggregate Reno throughput and link utilization at link 1. | 88 |
| 6.2 | Experiment 6.2: Aggregate Reno throughput and link utilization at link 2. | 89 |
| 6.3 | Experiment 6.3: Queue sizes at link 1 and link 2. The system shifts between the two equilibria with disjoint active constraint sets. | 90 |

Chapter 1

Introduction

...he intends only his own gain, and he is in this, as in many other cases, led by an **invisible hand** to promote an end which was no part of his intention.

— from “*The Wealth of Nations*” by Adam Smith

Adam Smith is generally regarded as the father of economics. His “invisible hand” idea has been widely adopted in the modern era to refer to a process in which the outcome to be explained is produced in a decentralized way (coordinated by the invisible hand), with no explicit agreements among the acting agents. The classic setting of the theory has a few basic assumptions. One fundamental yet frequently ignored assumption is the homogeneity of prices. In other words, the invisible hand (price signal) observed by all agents is assumed to be exactly the same. In this thesis, by studying a concrete engineering system (network congestion control) which has been modelled using market-based theory, we discover some consequences of relaxing the price homogeneity assumption. It is demonstrated that both the results and the mathematical techniques used to derive them are very different from the traditional theory and our systematic study provides predictions that are verified by experiments. Throughout this thesis, we mainly focus on network congestion control, whose basic concepts and results are introduced in this chapter. However, the discovery here can have potential impacts on related problems in both economics and mathematics that will be discussed later.

1.1 Network Congestion Control

Congestion control has long been recognized as a very important component of network regulation, in both traditional transportation networks and modern communication networks. It is of particular importance for a giant and heterogeneous network like the Internet, where congestion can potentially lead to huge performance degradation.

Congestion occurs when the aggregate demand for a certain resource (e.g., link bandwidth) exceeds its supply. The effects, using the Internet as an example, include long transfer delay, high packet loss, frequent packet retransmissions, and even possible congestion collapse [32], where network links are fully utilized but the throughput which an application obtains is close to zero. This indeed happened historically. In October 1986, the Internet had one of its first congestion collapses. A link between two terminals in the Lawrence Berkeley Laboratory and the University of California at Berkeley, which were about 400 yards apart had its throughput drop from 32 Kbps to 40 bps (about a factor of 1000) ! After a series of similar events, people began to build in the congestion control component into TCP (Transmission Control Protocol), which has been widely regarded as a great success and a significant contribution to the triumph of the Internet.

The basic idea behind TCP congestion control is straightforward. Data sources regulate their sending rates according to feedback signals (packet loss rate for current TCP) from the network. These signals are generated by links based on their utilization and hence carry information about the congestion level inside the network. It is worthwhile to note that data sources only need to know the aggregate congestion signals in their paths, and links update their congestion signals only based on aggregate traffic going through them. In other words, it is a distributed system and only local information is used.

Borrowing concepts from economics and tools from optimization and control theory, significant progress toward understanding network congestion control has been made in the last decade [39, 46, 55, 84, 41, 44]. In short, we view the rate (primal variable) adaptation at sources and the congestion signal (dual variable) adjustment at links as distributed primal/dual algorithms to solve a convex optimization problem. When the whole network reaches its equilibrium, the flow rates maximize aggregate utility while the congestion sig-

nals, which can be interpreted as market-clearing prices for those resources, solve its dual problem. For more mathematical details, see section 2.1.3.

This duality theory has its impact in many ways. First, it helps explain previous empirical observations. Second, by showing the global consequence of collective local algorithms, it sets up a framework under which properties like efficiency, and fairness can be rigorously discussed. Third, with its elegant and simple structure, it allows us to discover previously unknown behaviors of end-to-end congestion control for networks with arbitrary topology [74]. Finally, it can help design new protocols that scale better with bandwidth [35, 80].

1.2 A Motivating Example

The key step in setting up the duality theory for congestion control system is to view congestion signals associated with links as Lagrange multipliers for the corresponding capacity constraints introduced by links. This, however, directly implies the assumption of homogeneity of congestion signals as there is only one set of dual variables in the duality theory. This implicit assumption is clearly violated whenever there are data sources that use different congestion signals. We now argue that is an important case.

The currently deployed TCP implementation, referred to as TCP Reno¹ in this thesis, uses packet loss as its congestion signal to dynamically adapt transmission rate, or equivalently, congestion window size. It has worked remarkably well in the past, but its limitations in wireless networks and in networks with a large bandwidth-delay product have been well-known and have motivated various proposals that use different congestion signals. For example, schemes that use queueing delay include the early proposals CARD [33], DUAL [79] and Vegas [12], and the recent proposal FAST [35, 80]. Solutions that use one-bit congestion signals include ECN [59, 83, 43], and those that use multi-bit feedback include XCP [37, 45], MaxNet [82, 81], and RCP [23]. Indeed, the Linux operating system already allows users to choose from a variety of congestion control algorithms since kernel version

¹All our experiments and simulations use NewReno with SACK. These are enhanced versions of the original Tahoe and Reno, but we refer to them generically as TCP Reno.

2.6.13 [1]. Clearly, going forward, our network will become a more heterogeneous one in which protocols that react to different congestion signals interact. Yet, our understanding of such a network is rudimentary at best. With the exception of a few limited analysis on very simple topologies [54, 42, 44, 27], existing literature generally assumes that all sources are homogeneous in that, even though they may control their rates using different algorithms, they all adapt to the same type of congestion signals.

In the remainder of this section, we will experimentally show that the current duality theory cannot explain even some of the most basic behaviors of networks with heterogeneous protocols. This example motivates the study in this thesis.

TCP Reno and FAST are used in the experiments, and we give a brief introduction to them here. TCP Reno controls its rate based on end-to-end packet loss probability. At its equilibrium point, Reno is characterized by the following equation [44]:

$$q_i^r = \frac{2}{2 + (x_i^r)^2 T_i^2} \quad (1.1)$$

where q_i^r is the end-to-end loss probability. Here, we assume that the round-trip time T_i for Reno flow i is a constant.

FAST [80] is a high speed TCP variant that uses delay as its main control signal. Every 20ms, a FAST flow adjusts its congestion window W according to

$$W \leftarrow \frac{baseRTT}{RTT} W + \alpha \quad (1.2)$$

where RTT is the current round trip time and baseRTT is the minimum RTT that has been observed. At equilibrium, each FAST flow i achieves a throughput

$$x_i^* = \frac{\alpha}{q_i^*} \quad (1.3)$$

where q_i^* is the equilibrium queueing delay observed by flow i . Hence, α is the number of packets that each FAST flow maintains in the bottlenecks along its path.

We set up a Dummynet testbed [61] with seven Linux servers as senders and receivers and three BSD servers to emulate software routers; see Figure 1.1. The Linux senders and




Figure 1.1: Dummynet setup for Experiments 1.1 and 1.2.

receivers run TCP Reno or FAST. The three emulated routers run FreeBSD 5.2.1. Each testbed machine has dual Xeon 2.66GHz processors, 2GB of main memory, and dual on-board Intel PRO/1000 gigabit Ethernet interfaces. The test machines are interconnected through a Cisco 3750 gigabit switch. The network is fully configurable, and the link delay and capacity can be modified on the emulated router. The queueing discipline is Droptail. We have programmed the Dummynet router to capture various state variables for computing queue trajectories, loss, and bandwidth utilization. The sender and receiver hosts have been instrumented to monitor TCP state variables. We use a 2.4.22 modified FAST kernel. In order to minimize host limitations and accommodate large bursts, we have increased the Linux transmission queue length to 5000 and ring buffer to 4096. Iperf is used to generate TCP traffic for each protocol.

We make the following remarks before reporting experiments in detail:

- We have modified the FAST implementation so that it does not halve its window after a loss. Therefore it only reacts to queueing delay [76].
- The standard 1500-byte MTU (Maximum Transmission Unit) is used. Thus, 100 Mbps = 8.33 pkts/ms.
- All the queue sizes reported below are exponential moving averages of instantaneous queue trajectories. Averaging does not affect the equilibrium value, which is of primary interest here. Note, however, that even though the averaged trajectory may not

reach buffer capacity, the instantaneous trajectory often does.

The topology of the network is shown in Figure 1.2. Links 1 and 3 (which correspond to the outgoing links of routers 1 and 3) are each configured with 110 Mbps capacity, 50 ms one-way propagation delay and a buffer of 800 packets. Link 2 (router 2) has a capacity of 150 Mbps with 10 ms one-way propagation delay and a buffer size of 150 packets. There are 8 Reno flows on path 3 utilizing all the three links, with one-way propagation delay of 110 ms. There are two FAST flows on each of paths 1 and 2. Both of them have one-way propagation delays of 60 ms. All FAST flows use a common parameter $\alpha = 50$ packets.




Figure 1.2: Multiple equilibria scenario.

Experiment 1.1²: multiple equilibrium points

The goal of this experiment is to demonstrate multiple (two) equilibrium points of the system. It directly contradicts the uniqueness of equilibrium predicted by the duality theory (the equilibrium is the solution of the utility maximization problem, which is a strictly convex program and hence admits a unique solution).

Two sets of tests have been carried out with different starting times for Reno and FAST flows. The intuition is that if FAST flows start first, link 2 will be saturated and links 1 and 3 will not. Since the buffer size for link 2 is small, when Reno flows join, they will experience so many losses that links 1 and 3 will remain unsaturated. This corresponds to an equilibrium with a bottleneck link set consisting of link 2 only. If Reno starts first, on the other hand, links 1 and 3 are saturated while link 2 is not because link 2 has a higher capacity. Since the buffer size at links 1 and 3 is large, they can generate enough queueing delay to squeeze FAST flows when they join and keep link 2 unsaturated. This

²Throughout this thesis, packet level simulation and experiment are called "Experiment" while numerical or theoretical example are called "Example". They are numbered separately.

corresponds to an equilibrium whose bottleneck link set consists of links 1 and 3. We repeat the experiments 30 times for both scenarios and obtain the following results.

The average aggregate rates and the standard deviation over the 30 experiments of all the flows on each of paths 1, 2, and 3 are shown in Table 1.1 for both starting orders. Since

| | Path 1 (FAST) | Path 2 (FAST) | Path 3 (Reno) |
|------------------|------------------|------------------|------------------|
| FAST start first | (52.0, 2.0) Mbps | (61.1, 3.3) Mbps | (26.6, 4.8) Mbps |
| Reno start first | (13.3, 0.8) Mbps | (13.4, 0.8) Mbps | (92.7, 0.7) Mbps |

Table 1.1: Average aggregate rates and their standard deviations of all flows on paths 1, 2, 3.

the difference in the aggregate rates for each path is far more than the standard deviation, it is clear that the network reaches very different equilibria depending on which flows start first. This is further confirmed by queue and throughput measurements shown in Figure 1.3 for link 1 and in Figure 1.4 for link 2 for one of the 30 experiments. The results for link 3 are similar to those for link 1 and are omitted.

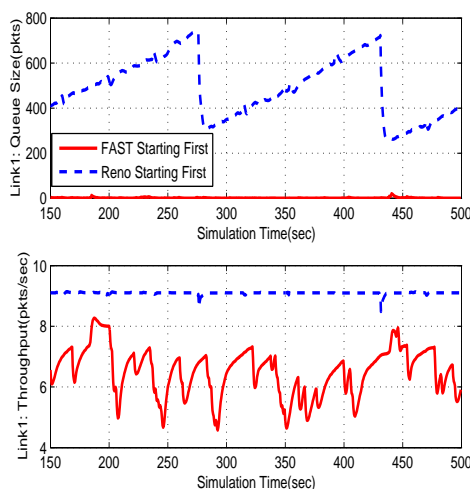


Figure 1.3: Experiment 1.1: queue size and aggregate throughput at link 1.

These figures show that when FAST flows start first, the link 2 queue remains nonzero while the queue of link 1 (and hence the link 3 queue) remains empty throughout the experiment, as expected. As a consequence, the aggregate throughput at link 2 is close to capacity while that at link 1 remains low for most of the time. When Reno flows start first,




Figure 1.4: Experiment 1.1: queue size and aggregate throughput at link 2.

the queue and throughput behaviors are exactly the opposite. For more detailed simulations on multiple equilibria, refer to Appendix 6.1.

Experiment 1.2: unique equilibrium point

To make sure that the above behavior is indeed caused by the existence of multiple protocols rather than by different flow arrival patterns, we repeat the experiment with the same network setup, but using all Reno or all FAST flows. When we use FAST flows along the long path, the parameter α is set to 30. The average throughput results are summarized in tables 1.2 and 1.3. They confirm that the network admits a unique equilibrium when a single protocol is used, regardless of flow arrival patterns.

| | Path 1 | Path 2 | Path 3 |
|-------------------------|------------------|------------------|------------------|
| Short flows start first | (47.7, 1.3) Mbps | (70.1, 1.8) Mbps | (13.4, 1.0) Mbps |
| Long flow starts first | (40.7, 1.5) Mbps | (64.9, 2.0) Mbps | (21.4, 1.2) Mbps |

Table 1.2: Average aggregate rates and their standard deviations of all flows on paths 1, 2, 3 (All flows are Reno).

| | Path 1 | Path 2 | Path 3 |
|-------------------------|------------------|------------------|------------------|
| Short flows start first | (47.2, 1.1) Mbps | (72.3, 1.6) Mbps | (15.6, 1.2) Mbps |
| Long flow starts first | (46.8, 1.3) Mbps | (72.0, 1.7) Mbps | (16.3, 1.0) Mbps |

Table 1.3: Average aggregate rates and their standard deviations of all flows on paths 1, 2, 3 (All flows are FAST).

1.3 Contributions of This Thesis

In this thesis, we first identify homogeneity of price as an implicit yet fundamental assumption in the current duality theory on price based flow control and explicitly show that, once this assumption is relaxed, the network can exhibit phenomena that the current theory fails to explain. We illustrate this through analysis, numerical examples, simulations, and experiments.

We make two main contributions. First, for networks with heterogeneous protocols, we examine all the basic issues of equilibrium including existence, uniqueness, optimality, and stability. These are described in turn in chapters 2 and 3. Second, in chapter 4, we propose an update in a slow timescale to help a network reach a unique and optimal equilibrium. Analysis and experiments are used to verify its correctness and effectiveness.

In chapter 2, we focus on existence and uniqueness of equilibrium. We prove the existence of equilibrium in general multi-protocol networks under mild assumptions. We show that for almost all networks, the equilibria are locally unique, and finite and odd in number. They cannot all be locally stable unless the equilibrium is globally unique. Finally, we show that if the price mapping functions that relate different prices observed by the sources are similar, global uniqueness is guaranteed. The similarity requirement is quantified.

In chapter 3, optimality and stability of equilibrium are discussed. By identifying an optimization problem associated with *every* equilibrium, we show that all equilibria are Pareto efficient. We also provide an upper bound on efficiency loss due to pricing heterogeneity. On fairness, we show that intra-protocol fairness is still decided by utility maximization problem while inter-protocol fairness is the part over which we don't have control. However it is shown that we can achieve any desirable inter-protocol fairness by properly choosing protocol parameters. Various stability results are provided. In particular we prove

that, if the degree of pricing heterogeneity is small enough, the network equilibrium is not only unique but also locally stable.

In chapter 4, motivated by the fact that there is in general efficiency loss with heterogeneous protocols, we shift our focus from analysis to design by asking the question of how to steer the network to the globally optimal operating point. We propose a scheme to steer an arbitrary network to a unique equilibrium that maximizes the total utility, by updating in a slow timescale a linear parameter in sources' algorithms. The scheme uses only local information. In addition to analysis, we present numerical examples, simulations and experiments using TCP Reno and FAST to demonstrate the correctness and convergence of the scheme.

Finally, we conclude in chapter 5 by discussing open issues and questions raised by this thesis. In addition to natural extension of current results, since the problem sits at the nexus of engineering and economics and our study has a lot to do with some problems in the frontier of linear algebra, we point out a few "bigger" related issues in those fields.

Chapter 6 is an appendix that contains detailed simulations and proofs that are too lengthy to be included in the main text without sacrificing fluency. Readers can skip them if they are primarily interested in the big picture and main results. However, chapter 6 does include interesting and original contributions such as the proofs of Theorem 2.8 and 3.8, which may have impacts on related mathematics and economics problems as discussed at the end of chapter 5.

Chapter 2

Existence and Uniqueness

Before discussing any performance metric of the congestion control system, which is evaluated at equilibrium, one needs to determine fundamental properties of equilibrium, like existence and uniqueness, which are the focus of this chapter. We prove that equilibrium still exists, under mild conditions, despite the lack of an underlying concave optimization problem (Section 2.2). In contrast to the single-protocol case, even when the routing matrix has full row rank, there can be uncountably many equilibria (Example 2.1 in Section 4.1) and the set of bottleneck links can be non-unique (Example 2.2 in Section 4.1). However, we prove that almost all networks have a finite number of equilibria and they are necessarily locally unique (Section 2.3.2). The number of equilibria is always odd, though can be more than one (Section 2.3.2). Moreover, these equilibria cannot all be locally stable unless the equilibrium is globally unique (Section 2.3.2). Finally, we provide two sufficient conditions for global uniqueness of network equilibrium (Sections 2.3.5 and 2.3.4). The first condition implies that if the price mapping functions that map link prices to effective prices observed by the sources do not differ too much, then global uniqueness is guaranteed. The second condition generalizes the full-rank condition on routing matrix for global uniqueness from single-protocol networks to multi-protocol networks. Throughout the section, we provide numerical examples and simulations to illustrate equilibrium properties or how theorems can be applied.

2.1 Model and Notations

2.1.1 Notations

A network consists of a set of L links, indexed by $l = 1, \dots, L$, with finite capacities c_l . We sometimes abuse notation and use L to denote both the number of links and the set $L = \{1, \dots, L\}$ of links. Each link has a price p_l as its congestion measure. There are J different protocols indexed by superscript j , and N^j sources using protocol j , indexed by (j, i) where $j = 1, \dots, J$ and $i = 1, \dots, N^j$. The total number of sources is $N := \sum_j N^j$.

The $L \times N^j$ routing matrix R^j for type j sources is defined by $R_{li}^j = 1$ if source (j, i) uses link l , and 0 otherwise. The overall routing matrix is denoted by

$$R = \begin{bmatrix} R^1 & R^2 & \dots & R^J \end{bmatrix}$$

Even though different classes of sources react to different prices, e.g., Reno to packet loss probability and Vegas/FAST to queueing delay, the prices are related. We model this relationship through a price mapping function that maps a common price (e.g., queue length) at a link to different prices (e.g., loss probability and queueing delay) observed by different sources. Formally, every link l has a price p_l . A type j source reacts to the "effective price" $m_l^j(p_l)$ in its path, where m_l^j is a price mapping function, which can depend on both the link and the protocol type. The exact form of m_l^j depends on the AQM algorithm used at the link. One can also take the price p_l^j used by one of the protocols, e.g., queueing delay, as the common price p_l . In this case the corresponding price mapping function is the identity function, $m_l^j(p_l) = p_l$. Taking a link with RED as an example and using delay as the common price p_l , the price mapping function m_l , which relates loss and delay, can now be explicitly expressed as:

$$p_l^r = m_l(p_l) = \begin{cases} 0 & p_l \leq \frac{b}{c_l} \\ \frac{1}{K} \frac{p_l c_l - b}{\bar{b} - b} & \frac{b}{c_l} \leq p_l \leq \frac{\bar{b}}{c_l} \\ \frac{1}{K} & p_l \geq \frac{\bar{b}}{c_l} \end{cases} \quad (2.1)$$

where \underline{b} , \bar{b} and K are RED parameters [72].

Let $m^j(p) = (m_l^j(p_l), l = 1, \dots, L)$ and $m(p) = (m^j(p_l), j = 1, \dots, J)$. The aggregate price for source (j, i) is defined as

$$q_i^j = \sum_l R_{li}^j m_l^j(p_l) \quad (2.2)$$

Let $q^j = (q_i^j, i = 1, \dots, N^j)$ and $q = (q^j, j = 1, \dots, J)$ be vectors of aggregate prices. Then $q^j = (R^j)^T m^j(p)$ and $q = R^T m(p)$.

Let x^j be a vector with the rate x_i^j of source (j, i) as its i th entry, and x be the vector of x^j :

$$x = \left[(x^1)^T, (x^2)^T, \dots, (x^J)^T \right]^T$$

Source (j, i) has a utility function $U_i^j(x_i^j)$ that is strictly concave and increasing in its rate x_i^j . Let $U = (U_i^j, i = 1, \dots, N^j, j = 1, \dots, J)$.

In general, if z_k are defined, then z denotes the (column) vector $z = (z_k, \forall k)$. Other notations will be introduced later when they are encountered. We call (c, m, R, U) a *network*.

2.1.2 Network equilibrium

A network is in equilibrium, or the link prices p and source rates x are in equilibrium, when each source (j, i) maximizes its net benefit (utility minus bandwidth cost), and the demand for and supply of bandwidth at each bottleneck link are balanced. Formally, a network equilibrium is defined as follows.

Given any prices p , we assume in this dissertation that the source rates x_i^j are uniquely determined by

$$x_i^j(q_i^j) = \left[(U_i^j)'^{-1}(q_i^j) \right]^+$$

where $(U_i^j)'$ is the derivative of U_i^j , and $(U_i^j)'^{-1}$ is its inverse, which exists since U_i^j is

strictly concave. Here $[z]^+ = \max\{z, 0\}$. This implies that the source rates x_i^j uniquely solve

$$\max_{z \geq 0} U_i^j(z) - zq_i^j$$

As we will see, under the assumptions in this thesis, $(U_i^j)^{\prime-1}(q_i^j) > 0$ for all the prices p that we consider, and hence we can ignore the projection $[\cdot]^+$ and assume without loss of generality that

$$x_i^j(q_i^j) = (U_i^j)^{\prime-1}(q_i^j) \quad (2.3)$$

As usual, we use $x^j(q^j) = (x_i^j(q_i^j), i = 1, \dots, N^j)$ and $x(q) = (x^j(q^j), j = 1, \dots, J)$ to denote the vector-valued functions composed of x_i^j . Since $q = R^T m(p)$, we often abuse notation and write $x_i^j(p), x^j(p), x(p)$.

Define the aggregate source rates $y(p) = (y_l(p), l = 1, \dots, L)$ at links l as:

$$y^j(p) = R^j x^j(p), \quad y(p) = R x(p) \quad (2.4)$$

In equilibrium, the aggregate rate at each link is no more than the link capacity, and they are equal if the link price is strictly positive. Formally, we call p an *equilibrium price*, a *network equilibrium*, or just an *equilibrium* if it satisfies (from (2.2)–(2.4))

$$P(y(p) - c) = 0, \quad y(p) \leq c, \quad p \geq 0 \quad (2.5)$$

where $P := \text{diag}(p_l)$ is a diagonal matrix. The goal of this chapter is to study the existence and uniqueness properties of network equilibrium specified by (2.2)–(2.5). Let E be the equilibrium set:

$$E = \{p \in \mathfrak{R}_+^L \mid P(y(p) - c) = 0, \quad y(p) \leq c\} \quad (2.6)$$

For future use, we now define an active constraint set and the Jacobian for links that

are actively constrained. Fix an equilibrium price $p^* \in E$. Let the *active constraint set* $\hat{L} = \hat{L}(p^*) \subseteq L$ (with respect to p^*) be the set of links l at which $p_l^* > 0$. Consider the reduced system that consists only of links in \hat{L} , and denote all variables in the reduced system by \hat{c} , \hat{p} , \hat{y} , etc. Then, since $y_l(p) = c_l$ for every $l \in \hat{L}$, we have $\hat{y}(\hat{p}) = \hat{c}$. Let the Jacobian for the reduced system be $\hat{J}(\hat{p}) = \partial \hat{y}(\hat{p}) / \partial \hat{p}$. Then

$$\hat{J}(\hat{p}) = \sum_j \hat{R}^j \frac{\partial x^j}{\partial \hat{q}^j}(\hat{p}) \left(\hat{R}^j \right)^T \frac{\partial \hat{m}^j}{\partial \hat{p}}(\hat{p}) \quad (2.7)$$

where

$$\frac{\partial x^j}{\partial \hat{q}^j} = \text{diag} \left(\left(\frac{\partial^2 U_i^j}{\partial (x_i^j)^2} \right)^{-1} \right) \quad (2.8)$$

$$\frac{\partial \hat{m}^j}{\partial \hat{p}} = \text{diag} \left(\frac{\partial \hat{m}_l^j}{\partial \hat{p}_l} \right) \quad (2.9)$$

and all the partial derivatives are evaluated at the generic point \hat{p} .

2.1.3 Current theory: $J = 1$

In this subsection, we briefly review the current theory for the case where there is only one protocol, i.e., $J = 1$, and explain why it cannot be directly applied to the case of heterogeneous protocols.

When all sources react to the same price, then the equilibrium described by (2.2)–(2.5) is the unique solution of the following utility maximization problem defined in [39]:

$$\max_{x \geq 0} \sum_i U_i(x_i) \quad (2.10)$$

$$\text{subject to } Rx \leq c \quad (2.11)$$

and its Lagrangian dual [46]:

$$\min_{p \geq 0} \sum_i \max_{x_i \geq 0} \left(U_i(x_i) - x_i \sum_l R_{li} p_l \right) + \sum_l c_l p_l \quad (2.12)$$

where we have omitted the superscript $j = 1$. The continuity and strict concavity of U_i plus the compactness of the feasible set guarantee the existence and uniqueness of the optimal solution of (2.10)–(2.11).

The basic idea in relating the utility maximization problem (2.10)–(2.11) to the equilibrium equations (2.2)–(2.5) is to examine the dual of the utility maximization problem, and interpret the effective price $m_l(p_l)$ as a Lagrange multiplier associated with each link capacity constraint (see, e.g., [46, 55, 44]). As long as $m_l(p_l) \geq 0$ and $m_l(0) = 0$, one can replace p_l in (2.5) by $m_l(p_l)$. The resulting equation together with (2.2)–(2.4) provides the necessary and sufficient condition for $x_i(p)$ and $m_l(p_l)$ to be the primal and dual optimal, respectively.

This approach breaks down when there are $J > 1$ types of prices because there cannot be more than one Lagrange multiplier at each link. In general, an equilibrium no longer maximizes aggregate utility, nor is it unique as we have already seen in experiments in section 1.2. However, as shown in the next section, existence of equilibrium is still guaranteed under the following assumptions:

A1: Utility functions U_i^j are strictly concave and increasing, and twice continuously differentiable in their domains. Price mapping functions m_l^j are continuously differentiable in their domains and strictly increasing with $m_l^j(0) = 0$.

A2: For any $\epsilon > 0$, there exists a number p_{\max} such that if $p_l > p_{\max}$ for link l , then

$$x_i^j(p) < \epsilon \text{ for all } (j, i) \text{ with } R_{li}^j = 1$$

These are mild assumptions. Concavity and monotonicity of utility functions are often assumed in network pricing for elastic traffic. Moreover, most TCP algorithms proposed or deployed turn out to have strictly concave increasing utility functions; see e.g. [44]. The assumption on m_l^j preserves the relative order of prices and maps zero price to zero effective price. Assumption A2 says that when p_l is high enough, every source going through link l has a rate less than ϵ , which is satisfied by all TCPs.

2.2 Existence

In this section, we prove the existence of network equilibrium. We start with a lemma that bounds the equilibrium prices.

Lemma 2.1. *Suppose A1 and A2 hold. Given a network (c, m, R, U) , there is a scalar p_{\max} that upper bounds any equilibrium price p , i.e., $p_l \leq p_{\max}$ for all l .*

Proof. Choose $\epsilon = \min_l c_l/N$, and let p_{\max} be the corresponding scalar in A2. Suppose that there exists an equilibrium price p and a link l , such that $p_l > p_{\max}$. A2 implies that the aggregate equilibrium rate at link l satisfies

$$\sum_j \sum_i R_{li}^j x_i^j(p) < N\epsilon = \min_l c_l$$

Therefore, we get a link with $p_l > 0$ but not fully utilized. It contradicts the equilibrium condition (2.5). \square

The following theorem asserts the existence of equilibrium for a multi-protocol network.

Theorem 2.2. *Suppose A1 and A2 hold. There exists an equilibrium price p^* for any network (c, m, R, U) .*

Proof. Let p_{\max} be the scalar upper bound in Lemma 2.1. For any $p \in [0, p_{\max}]^L$, define a vector function

$$F(p) := Rx(p) - c \tag{2.13}$$

For any link l , let

$$p_{-l} := (p_1, \dots, p_{l-1}, p_{l+1}, \dots, p_L)^T$$

Then we may write $F(p)$ as $F(p_l, p_{-l})$. Define function h_l as

$$h_l(p_l, p_{-l}) := -F_l^2(p_l, p_{-l}) \tag{2.14}$$

We claim that $h_l(p_l, p_{-l})$ is a quasi-concave function in p_l for any fixed p_{-l} . By the definition of quasi-concavity in [56], we only need to check that the set

$$A_l := \{ p_l \mid h_l(p_l, p_{-l}) \geq a \}$$

is convex for all $a \in \mathfrak{R}$. If $a > 0$, clearly $A_l = \emptyset$ by (2.14). When $a \leq 0$, the set A_l can be rewritten as

$$A_l = \left\{ p_l \mid -\sqrt{|a|} \leq F_l(p_l, p_{-l}) \leq \sqrt{|a|} \right\}$$

Since $F_l(p_l, p_{-l})$ is a nonincreasing function in p_l for any fixed p_{-l} , the set A_l is convex. Therefore $h_l(p_l, p_{-l})$ is quasi-concave in p_l .

Since $[0, p_{\max}]$ is a nonempty compact convex set, by the theorem of Nash [56], the quasi-concavity of $h_l(p_l, p_{-l})$ guarantees that there exists a $p^* \in [0, p_{\max}]^L$ such that for all $l \in \{1, 2, \dots, L\}$

$$p_l^* = \arg \max_{p_l \in [0, p_{\max}]} h_l(p_l, p_{-l}^*)$$

We now argue that, for all l , either 1) $F_l(p^*) = 0$, or 2) $F_l(p^*) < 0$ and we can take $p_l^* = 0$. These conditions imply (2.5), and hence p^* is an equilibrium price.

Case 1: $F_l(0, p_{-l}^*) > 0$. Since U_i^j is strictly concave, $F_l(p_l, p_{-l}^*)$ is nonincreasing¹ in $[0, p_{\max}]$. Moreover, the proof of Lemma 2.1 shows that $F_l(p_{\max}, p_{-l}^*) < 0$. Therefore, there exists a point p_l^* in $[0, p_{\max}]$ where $F_l(p_l, p_{-l}^*) = 0$. This p_l^* maximizes $h_l(p_l, p_{-l}^*)$.

Case 2: $F_l(0, p_{-l}^*) \leq 0$. Since $F_l(p_l, p_{-l}^*)$ is a nonincreasing function in p_l , we have that

$$F_l(p_l, p_{-l}^*) \leq 0 \text{ for all } p_l \in [0, p_{\max}]$$

If $-c_l < F_l(0, p_{-l}^*) \leq 0$, then $F_l(p_l, p_{-l}^*)$ and $h_l(p_l, p_{-l}^*)$ are strictly decreasing in p_l and

¹ $F_l(p_l, p_{-l}^*)$ is strictly decreasing unless some $x_i(p)$ becomes zero.

hence

$$p_l^* = \arg \max_{p_l \in [0, p_{\max}]} h_l(p_l, p_{-l}^*) = 0$$

Otherwise we have $F_l(0, p_{-l}^*) = -c_l$ from (2.13). In this situation, all x_i^j going through link l are zero, and hence we can set $p_l^* = 0$ without affecting any other prices. More precisely, a (possibly) new price vector \tilde{p} with $\tilde{p}_l = 0$ and $\tilde{p}_k = p_k^*$ for $k \neq l$ is also a Nash equilibrium that maximizes $h_k(p_k, \tilde{p}_{-k})$ for $k = 1, \dots, L$.

Thus we have proved that, for $l = 1, \dots, L$,

$$p_l^* F_l(p_l^*, p_{-l}^*) = 0, \quad F_l(p_l^*, p_{-l}^*) \leq 0, \quad p^* \geq 0$$

which is (2.5). □

2.3 Uniqueness

2.3.1 Multiple equilibria: examples

In a single-protocol network, if the routing matrix R has full row rank, then there is a unique active constraint set \hat{L} and a unique equilibrium price p associated with it. If R does not have full row rank, then equilibrium prices p may be nonunique but the equilibrium rates $x(p)$ are still unique since the utility functions are strictly concave, and the feasible set is convex.

In contrast, there can be multiple equilibrium prices associated with the same active constraint set (Example 2.1). Moreover, the active constraint set in a multi-protocol network can be nonunique even if R has full row rank (Example 2.2). Clearly, the equilibrium prices associated with different active constraint sets are different.

Example 2.1: unique active constraint set but uncountably many equilibria

In this example, we assume all of the sources use the same utility function

$$U_i^j(x_i^j) = -\frac{1}{2}(1 - x_i^j)^2 \quad (2.15)$$

Then the equilibrium rates x^j of type j sources are determined by the equilibrium prices p as

$$x^j(p) = \mathbf{1} - (R^j)^T m^j(p)$$

where $\mathbf{1}$ is a vector of appropriate dimension whose entries are all 1s. We use linear price mapping functions:

$$m^j(p) = K^j p$$

where K^j are $L \times L$ diagonal matrices. Then the equilibrium rate vector of type j sources can be expressed as

$$x^j(p) = \mathbf{1} - (R^j)^T K^j p$$

When only links with strictly positive equilibrium prices are included in the model, we have

$$y(p) = \sum_{j=1}^J R^j x^j(p) = c$$

Substituting in $x^j(p)$ yields

$$\sum_{j=1}^J R^j (R^j)^T K^j p = \sum_{j=1}^J R^j \mathbf{1} - c$$

which is a linear equation in p for given R^j , K^j , and c . It has a unique solution if the

determinant is nonzero, but has no or multiple solutions if

$$\det \left(\sum_{j=1}^J R^j (R^j)^T K^j \right) = 0$$

When $J = 1$, i.e., there is only one protocol, and R^1 has full row rank, $\det(R^1 (R^1)^T K^1) > 0$ since both $R^1 (R^1)^T$ and K^1 are positive definite. In this case, there is a unique equilibrium price vector. When $J = 2$, there are networks whose determinants are zero that have uncountably many equilibria. See Appendix 6.4 for an example where R does not have full row rank. We provide here an example with $J = 3$ where R still has full row rank.

The network is shown in Figure 2.1 with three unit-capacity links, $c_l = 1$.




Figure 2.1: Example 2.1: uncountably many equilibria.

There are three different protocols with the corresponding routing matrices

$$R^1 = I, \quad R^2 = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}^T, \quad R^3 = (1, 1, 1)^T$$

The linear price mapping functions are given by

$$K^1 = I, \quad K^2 = \text{diag}(5, 1, 5), \quad K^3 = \text{diag}(1, 3, 1)$$

It is easy to calculate that

$$\sum_{i=1}^3 R^i (R^i)^T K^i = \begin{bmatrix} 7 & 4 & 1 \\ 6 & 6 & 6 \\ 1 & 4 & 7 \end{bmatrix}$$

which has determinant 0. Using the utility function defined in (2.15), we can check that the following are equilibrium prices for any $\epsilon \in [0, 1/24]$:

$$p_1^1 = p_3^1 = 1/8 + \epsilon \quad p_2^1 = 1/4 - 2\epsilon$$

The corresponding equilibrium rates are

$$\begin{aligned} x_1^1 = x_3^1 &= 7/8 - \epsilon & x_2^1 &= 3/4 + 2\epsilon \\ x_1^2 = x_2^2 &= 1/8 - 3\epsilon & x_1^3 &= 4\epsilon \end{aligned}$$

All capacity constraints are tight with these rates. Since there is a one-link flow at every link, the active constraint set is unique and contains every link. Yet there are uncountably many equilibria.

Example 2.2: multiple active constraint sets each with a unique equilibrium

Consider the symmetric network in Figure 2.2 with 3 flows, which is the same topology used experiments in section 1.2.




Figure 2.2: Two equilibrium with different active bottleneck link sets.

There are two protocols in the network with the following routing matrices:

$$R^1 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad R^2 = (1, 1, 1)^T$$

Flows (1, 1) and (1, 2) have identical utility function U^1 and source rate x^1 , and flow (2, 1) has a utility function U^2 and source rate x^2 .

Links 1 and 3 both have capacity c_1 and price mapping functions $m_1^1(p) = m_3^1(p) = p$ and $m_1^2(p)$ for protocols 1 and 2, respectively. Link 2 has capacity c_2 and price mapping functions $m_2^1(p) = m_3^2(p) = p$ and $m_2^2(p)$.

Theorem 2.3. *Suppose assumption A1 holds. The network shown in Figure 4.4 has two equilibria provided:*

1. $c_1 < c_2 < 2c_1$;
2. For $j = 1, 2$, $(U^j)'(x^j) \rightarrow \bar{p}^j$ for some \bar{p}^j possibly ∞ , if and only if $x^j \rightarrow 0$;
3. For $l = 1, 2$, $m_l(p_l) \rightarrow \bar{p}^l$ as $p_l \rightarrow \bar{p}^l$, and satisfy

$$\begin{aligned} 2m_1((U^1)'(c_2 - c_1)) &< (U^2)'(2c_1 - c_2) \\ &< m_2((U^1)'(c_2 - c_1)) \end{aligned}$$

Proof: We first claim that, if $c_1 < c_2$ and $(U^2)'(2c_1 - c_2) > 2m_1^2((U^1)'(c_2 - c_1))$, then there is an equilibrium point where only links 1 and 3 are saturated and link 2 is not. In this case the equilibrium price for link 2 is $p_2 = 0$ and, by symmetry, those for links 1 and 3 are both p_1 . Such an equilibrium, if exists, is defined by the following equations:

$$\begin{aligned} (U^1)'(x^1) &= p_1 & (U^2)'(x^2) &= 2m_1(p_1) \\ x^1 + x^2 &= c_1 & 2x^1 + x^2 &< c_2 \end{aligned}$$

Eliminating x^2 and p_1 , the above equations are reduced to:

$$(U^2)'(c_1 - x^1) = 2m_1((U^1)'(x^1)) \quad (2.16)$$

$$x^1 < c_2 - c_1 \quad (2.17)$$

An equilibrium exists if and only if (2.16)–(2.17) has a nonnegative solution for x^1 . We now show that (2.16)–(2.17) indeed admits a unique solution $x^* > 0$ under the hypothesis of the proposition.

When $x^1 = 0$, we have

$$(U^2)'(c_1 - x^1) = (U^2)'(c_1) < \bar{p}^2 \leq 2\bar{p}^2 = 2m_1((U^1)'(0))$$

The inequality and the last equality have made multiple use of conditions 2 and 3 of the proposition. On the other hand, when $x^1 = c_2 - c_1$, we have $U_2'(2c_1 - c_2) > 2m_1(U_1'(c_2 - c_1))$ by condition 3. Since all functions here are continuous, $(U^j)'$ are strictly decreasing, and m_l are strictly increasing, there exists a unique $0 < x^* < c_2 - c_1$ such that $(U^2)'(c_1 - x^*) = 2m_1((U^1)'(x^*))$.

We next claim that, if $c_2 < 2c_1$ and $(U^2)'(2c_1 - c_2) < m_2((U^1)'(c_2 - c_1))$, then there is an equilibrium point where only link 2 is saturated and links 1 and 3 are not. In this case $p_1 = 0$, and the following equations determine such an equilibrium:

$$(U^1)'(x^1) = p_2 \quad (U^2)'(x^2) = m_2(p_2)$$

$$x^1 + x^2 < c_1 \quad 2x^1 + x^2 = c_2$$

Eliminating x^2 and p_2 , the equilibrium is specified by

$$(U^2)'(c_2 - 2x^1) = m_2((U^1)'(x^1)) \quad (2.18)$$

$$x^1 > c_2 - c_1 \quad (2.19)$$

When $x^1 = c_2 - c_1$, we have

$$(U^2)'(c_2 - 2x^1) = (U^2)'(2c_1 - c_2) < m_2((U^1)'(x^1))$$

by condition 3. When $x^1 = c_2/2$

$$(U^2)'(c_2 - 2x^1) = (U^2)'(0) = \bar{p}^2 > m_2((U^1)'(x^1))$$

where we have used conditions 2 and 3. Hence, again, there is a unique x^* that satisfies (2.18)–(2.19). Moreover, from (2.17) and (2.19), the two equilibria are distinct. \square

2.3.2 Regular networks

Examples 2.1 and 2.2 show that global uniqueness is generally not guaranteed in a multi-protocol network. We now show, however, that local uniqueness is basically a generic property of the equilibrium set. We present our main results on the structure of the equilibrium set here, providing conditions for the equilibrium points to be locally unique, finite and odd in number, and globally unique. Proofs of these results are provided in the next subsection.

Consider an equilibrium price $p^* \in E$. Recall the active constraint set \hat{L} defined by p^* . The equilibrium price \hat{p}^* for the links in \hat{L} is a solution of

$$\hat{y}(\hat{p}) = \hat{c} \tag{2.20}$$

By the inverse function theorem, the solution of (2.20), and hence the equilibrium price \hat{p}^* , is *locally unique* if the Jacobian matrix $\hat{J}(\hat{p}^*) = \partial \hat{y} / \partial \hat{p}$ is nonsingular at \hat{p}^* . We call a network (c, m, R, U) *regular* if all its equilibrium prices are locally unique.

The next result shows that almost all networks are regular, and that regular networks have finitely many equilibrium prices. This justifies restricting our attention to regular networks and allows us to further characterize the structure of equilibrium set by using index theorem.

Theorem 2.4. *Suppose assumptions A1 and A2 hold. Given any price mapping functions m , any routing matrix R and utility functions U ,*

1. *The set of link capacities c for which not all equilibrium prices are locally unique has Lebesgue measure zero in \mathfrak{R}_+^L .*
2. *The number of equilibria for a regular network (c, m, R, U) is finite.*

For the rest of this subsection, we narrow our attention to networks that satisfy an additional assumption:

A3: Every link l has a single-link flow (j, i) with $(U_i^j)'(c_l) > 0$.

Assumption A3 says that when the price of link l is small enough, the aggregate rate through it will exceed its capacity. This ensures that the active constraint set contains all links and facilitates the application of Poincare-Hopf theorem by avoiding equilibrium on the boundary (some $p_l = 0$).²

Since all the equilibria of a regular network have nonsingular Jacobian matrices, we can define the *index* $I(p)$ of $p \in E$ as

$$I(p) = \begin{cases} 1 & \text{if } \det(\mathbf{J}(p)) > 0 \\ -1 & \text{if } \det(\mathbf{J}(p)) < 0 \end{cases}$$

Then, we have a global characterization of equilibrium set stated as the next theorem.

²It is recently shown in [67] that A3 is not necessary and one can generalize Theorem 2.5 to

$$\sum_{p \in E} (-1)^{\hat{L}(p)} I(p) = 1$$

where $\hat{L}(p)$ is the number of links of the active constraint set associated with equilibrium p . Clearly, if $\hat{L}(p) = L$, it reduces to Theorem 2.5. This generalized theorem also allows us to conclude the number of equilibria is odd (and therefore existence) without A3. In this dissertation, although A3 is imposed for ease of presentation, all results can be viewed with respect to a fixed active constraint set with appropriate modifications. In particular, the global uniqueness results in section 2.3.4 directly apply without A3 since $\hat{\mathbf{J}}$ has the same structure as \mathbf{J} except with a smaller dimension.

Theorem 2.5. *Suppose assumptions A1–A3 hold. Given any regular network, we have*

$$\sum_{p \in E} I(p) = (-1)^L$$

where L is the number of links.

To help appreciate Theorem 2.5, we give its two important consequences.

Corollary 2.6. *Suppose assumptions A1–A3 hold. A regular network has an odd number of equilibria.*

Notice that Corollary 2.6 implies the existence of equilibrium. Although we proved this in section 2.2 in a more general setting, this simple corollary shows the power of Theorem 2.5.

The next result provides a condition for global uniqueness. We say that an equilibrium $p^* \in E$ is *locally stable* if the corresponding Jacobian matrix $\mathbf{J}(p^*)$ defined in (2.7) is stable, that is, every eigenvalue of $\mathbf{J}(p^*) = \partial y(p^*)/\partial p$ has negative real part. For justification of this definition, local stability of p^* implies that the gradient algorithm (2.23) later converges locally.

Corollary 2.7. *Suppose assumptions A1–A3 hold. The equilibrium of a regular network is globally unique if and only if every equilibrium point in E has an index $(-1)^L$. In particular, if all equilibria are locally stable, then E contains exactly one point.*

This result may seem surprising at first sight as it relates the local stability of an algorithm to the uniqueness property of a network. This is because both equilibrium and local stability are defined in terms of the function $y(p)$: an equilibrium p^* satisfies $y(p^*) = c$ and the local asymptotic stability of p^* is determined by $\partial y(p^*)/\partial p$. The connection between these two properties is made exact by the index theorem. An implication of this result is that if there are multiple equilibria, then no algorithm $\dot{p} = f(p(t))$, whose linearization around each equilibrium $p^* \in E$ satisfies $\partial f(p^*)/\partial p = \Lambda \partial y(p^*)/\partial p$, can be found to locally stabilize all of the equilibria.

Corollary 2.7 will be used in Section 2.3.4 to derive a sufficient condition on price mapping functions m for global uniqueness. We close this subsection with an example that illustrates Theorem 2.5 and Corollaries 2.6 and 2.7.

Example 2.3: illustration of Theorem 2.5 and Corollaries 2.6, 2.7

We revisit Example 2.1 with different utility functions. Recall that in Example 2.1, as ϵ varies from 0 to $1/24$, we trace out all equilibrium points. The components x_1^1 and $q_1^1 = p_1^1$ of these equilibrium points are shown by the (red) solid line in Figure 2.3. Other sources x_i^j and their effective end-to-end prices q_i^j also lie on similar straight lines. Since the network has uncountably many equilibrium points, it is not regular. To make it regular, suppose we change the utility functions of sources (j, i) to

$$U_i^j(x_i^j, \alpha_i^j) = \begin{cases} \beta_i^j (x_i^j)^{1-\alpha_i^j} / (1 - \alpha_i^j) & \text{if } \alpha_i^j \neq 1 \\ \beta_i^j \log x_i^j & \text{if } \alpha_i^j = 1 \end{cases}$$

with appropriately chosen positive constants α_i^j and β_i^j . These utility functions can be viewed as a weighted version of the widely used α -fairness utility functions proposed in [55].

The basic idea of how to choose α_i^j and β_i^j to generate only finitely many equilibrium points is as follows. First, we pick two points in the equilibrium set of Example 2.1, say, the points associated with $\epsilon = 0.01$ and $\epsilon = 0.04$. These choices of ϵ provide two distinct equilibrium points (q, x) and (\tilde{q}, \tilde{x}) . For instance, $(q_1^1, x_1^1) = (0.135, 0.865)$ corresponds to $\epsilon = 0.01$ and $(\tilde{q}_1^1, \tilde{x}_1^1) = (0.165, 0.835)$ corresponds to $\epsilon = 0.04$, as illustrated in Figure 2.3. Then, for each source (j, i) , find α_i^j and β_i^j such that (2.3) is satisfied by the two equilibrium points (q_i^j, x_i^j) and $(\tilde{q}_i^j, \tilde{x}_i^j)$ with the new utility functions. This is illustrated in Figure 2.3 where relation (2.3) with the new utility function is represented by the (blue) curve, and α_i^j , β_i^j are chosen so that the curve passes through the original equilibrium points (x_1^1, q_1^1) and (\tilde{q}, \tilde{x}) . More specifically, given two equilibrium points (q_i^j, x_i^j) and $(\tilde{q}_i^j, \tilde{x}_i^j)$, choose

$$\alpha_i^j = \frac{\log(q_i^j) - \log(\tilde{q}_i^j)}{\log(\tilde{x}_i^j) - \log(x_i^j)} \quad \beta_i^j = q_i^j (x_i^j)^{\alpha_i^j}$$




Figure 2.3: Example 2.3: construction of multiple isolated equilibria.

The resulting α_i^j and β_i^j for all flows (j, i) are shown in Table 2.1.

Table 2.1: Example 2.3: α_i^j and β_i^j .

| Flows | α_i^j | β_i^j |
|---------|--------------|-------------|
| x_1^1 | 5.6851 | 0.0592 |
| x_2^1 | 4.0285 | 0.0803 |
| x_3^1 | 5.6851 | 0.0592 |
| x_1^2 | 0.0322 | 0.8389 |
| x_2^2 | 0.0322 | 0.8389 |
| x_1^3 | 0.0963 | 0.7041 |

By construction, both $(p_1^1 = 0.135, p_2^1 = 0.230)$ and $(p_1^1 = 0.165, p_2^1 = 0.170)$ are network equilibria. By Corollary 2.6, there is at least one additional equilibrium. Numerical search indeed located a third equilibrium with $(p_1^1 = 0.142, p_2^1 = 0.206)$.

We further check the local stability of these three equilibria under the gradient algorithm (2.23) to be introduced in Section 2.3.3. The eigenvalues and index for each equilibrium are shown in Table 2.2. It turns out that the equilibrium $(p_1^1 = 0.142, p_2^1 = 0.206)$ is not stable and has index 1, while the other two are stable with index -1 . The dynamics of

Table 2.2: Example 2.3: stability and indices of equilibria.

| Equilibria (p_1, p_2, p_3) | Eigenvalues | Index |
|------------------------------|-------------------------|-------|
| $(0.135, 0.23, 0.135)$ | $-0.21, -17.43, -26.73$ | -1 |
| $(0.142, 0.206, 0.142)$ | $0.21, -12.32, -22.40$ | 1 |
| $(0.165, 0.17, 0.165)$ | $-12.41, -1.67, -0.67$ | -1 |

this network under the gradient algorithm is illustrated by a vector field. By symmetry, the equilibrium prices for the first and third link are always same. Therefore, we can draw the vector field restricted on the plane $p_1 = p_3$ to demonstrate the dynamics of the system. That is shown in Figure 2.4. The (red) dots represent the three equilibria. Note the equilibrium in the middle is a saddle point, and therefore unstable. The (red) arrows give the direction of this vector field. Individual trajectories are plotted with thin (blue) lines.




Figure 2.4: Example 2.3: vector field of (p_1, p_2) .

2.3.3 Proofs

In this subsection we provide proofs for the results in Section 2.3.2.

Proof of Theorem 2.4. The main mathematical tool used in our proof is Sard's Theorem [18, 70], of which we quote a version here that is tailored to our problem. Let G be an open subset of \mathfrak{R}_+^L and let F be a continuously differentiable function from G to \mathfrak{R}_+^L . A point $y \in G$ is a *critical point* of F if the Jacobian matrix $\partial F / \partial y$ of F at y is singular. A point $z \in \mathfrak{R}_+^L$ is a *critical value* of F if there is a critical point $y \in G$ with $z = F(y)$. A point in \mathfrak{R}_+^L is a *regular value* of F if it is not a critical value.

Sard's theorem. If $F : G \rightarrow \mathfrak{R}_+^L$ is continuously differentiable on the open subset $G \subseteq \mathfrak{R}_+^L$, then the set of critical values of F has Lebesgue measure zero in \mathfrak{R}_+^L .

Fix a routing matrix R and utility functions U . There are at most $2^L - 1$ different active

constraint sets. Let $\hat{L} \subseteq L$ be such a combination with \hat{L} links. Consider the set of all possible link capacities $c = (c_l, l \in L)$ under which the active constraint set is \hat{L} , i.e., with such a capacity vector c , an equilibrium price p has $p_l > 0$ if $l \in \hat{L}$ and $p_l = 0$ otherwise. Fix such an equilibrium point p^* . Again let \hat{p} denote the price vector only for links in \hat{L} . Then \hat{p}^* is not locally unique if the function $\hat{y} : \mathfrak{R}_+^{\hat{L}} \rightarrow \mathfrak{R}_+^{\hat{L}}$ defined by $\hat{y}(\hat{p}) = \hat{R}x(\hat{p})$ has a singular Jacobian matrix $\partial\hat{y}/\partial\hat{p}$ at \hat{p}^* , i.e., if \hat{p}^* is a critical point of \hat{y} . The set of such capacity vectors $\hat{c} \in \mathfrak{R}_+^{\hat{L}}$ under which all links in \hat{L} have active constraints in equilibrium satisfy

$$\hat{y}(\hat{p}^*) = \hat{c}$$

and hence are critical values of \hat{y} . Since \hat{y} is continuously differentiable by assumption A1, we can apply Sard's theorem and conclude that the set of such capacity vectors \hat{c} has zero Lebesgue measure in $\mathfrak{R}_+^{\hat{L}}$. The extension to \mathfrak{R}_+^L for all link capacities clearly also has zero Lebesgue measure in \mathfrak{R}_+^L .

Since we only have a finite number of different active constraint sets, the union of link capacity vectors that give rise to locally nonunique equilibria still has zero Lebesgue measure. This proves the first part of the theorem.

The equilibrium set E defined in (2.6) is closed because $y(p)$ is continuous, and is bounded by Lemma 2.1. Hence E is compact. Since (c, m, R, U) is a regular network, every $p \in E$ is locally unique, i.e., for each $p \in E$ we can find an open neighborhood such that it is the only equilibrium in that open set. The union of these open sets forms a cover for set E . Since E is compact, it admits a finite subcover [48], i.e., E can be covered by a finite number of open sets each containing a single equilibrium. Hence, the number of equilibria is finite. \square

Proof of Theorem 2.5. By assumption A3, we can always find $p_{\min} > 0$ such that for any price p and link l with $p_l < p_{\min}$, we have

$$\sum_j \sum_i R_{li}^j x_i^j(p) > c_l$$

Let $G := [p_{\min}, p_{\max}]^L$ where p_{\max} is defined in Lemma 2.1. Clearly, all equilibria are in the set G . To prove our result, we will invoke a version of the Poincare-Hopf index theorem tailored to our problem [78, 53].

Poincare-Hopf index theorem. Let D be an open subset of \Re and $v : D^L \rightarrow \Re^L$ be a smooth vector field, with nonsingular Jacobian matrix $\partial v / \partial p$ at every equilibrium. If there is a $G \subseteq D^L$ such that every trajectory moves inward of region G , then the sum of the indices of the equilibria in G is $(-1)^L$.

Gradient project algorithm. To construct the vector field v required by the index theorem, let $D^L = G$ and consider the following gradient algorithm from G to G proposed in [46]. The prices are updated at time t according to

$$\dot{p}(t) = \Lambda (Rx(t) - c) \quad (2.21)$$

where $\Lambda > 0$ is an $L \times L$ diagonal matrix with all elements being positive. A source updates its rate based on the end-to-end price

$$x(t) = x(p(t)) \quad (2.22)$$

A consequence of assumption A3 is that $p(t) \geq p_{\min} > 0$ for all t under the gradient algorithm (2.21)–(2.22). This guarantees a unique active constraint set that is L . Hence the equilibrium set E defined in (2.6) is equivalent to $E = \{p \in \Re_+^L \mid y(p) - c = 0\}$.

Combining (2.21)–(2.22) with $y(p(t)) = Rx(t)$ yields the required vector field v :

$$\dot{p}(t) = \Lambda(y(p(t)) - c) =: v(p(t)) \quad (2.23)$$

whose Jacobian matrix is:

$$\frac{\partial v}{\partial p}(p) = \Lambda \mathbf{J}(p) = \Lambda \frac{\partial y}{\partial p}(p) \quad (2.24)$$

where $\mathbf{J}(p)$ is given by (2.7). Clearly, p^* is an equilibrium point of v , i.e., $v(p^*) = 0$, if and

only if p^* is a network equilibrium, i.e., $p^* \in E$. Since the network (c, m, R, U) is regular, $\mathbf{J}(p)$ is nonsingular at every network equilibrium $p^* \in E \subset G$. Since Λ is a positive diagonal matrix, $\partial v(p)/\partial p$ is also nonsingular by (2.24) at all its equilibrium points p in G , as the index theorem requires.

Consider any point p on the boundary of G . For any l , we have one of two cases:

1. If $p_l(t) = p_{\max}$, link l will be underutilized, $y_l(p(t)) < c_l$, and $\dot{p}_l < 0$ according to (2.23).
2. If $p_l(t) = p_{\min}$, the aggregate rate at link l will exceed c_l , $y_l(p(t)) > c_l$, and $\dot{p}_l > 0$ according to (2.23).

Therefore, every point p on the boundary of G will move inward and our result directly follows from the Poincare-Hopf index theorem. \square

Proof of Corollary 2.6. Since both $I(p)$ and $(-1)^L$ are odd, the number of terms in the summation in Theorem 2.5 must be odd. \square

Proof of Corollary 2.7. The first claim of the theorem directly follows from Theorem 2.5. We now claim that an equilibrium $p^* \in E$ which is locally stable has an index $I(p^*)$ of $(-1)^L$. To prove the claim, consider a locally stable equilibrium price p^* . All the eigenvalues of $\mathbf{J}(p^*)$ have negative real parts. Moreover, since $\mathbf{J}(p^*)$ has real entries, complex eigenvalues come in conjugate pairs. The determinant of $\mathbf{J}(p^*)$ is the product of all its eigenvalues. If there are k conjugate pairs of complex eigenvalues and $L - 2k$ real eigenvalues, the product of all eigenvalues has the same sign as $(-1)^{L-2k}$, which has the same sign as $(-1)^L$. Hence the index of a locally stable equilibrium is $(-1)^L$. \square

2.3.4 Global uniqueness: mapping functions $m(p)$

In this and the next subsections, we provide sufficient conditions on the structure of the network for global uniqueness. We also provide some important special cases in section 2.3.6 where global uniqueness is guaranteed. In this subsection, we reveal that, if the price mapping functions m_l^j are similar, then the equilibrium of a regular network is globally unique.

2.3.4.1 General result

To state the result concisely, we need the notion of permutation. We call a vector $\sigma = (\sigma_1, \dots, \sigma_L)$ a *permutation* if each σ_l is distinct and takes value in $\{1, \dots, L\}$. Treating σ as a mapping $\sigma : \{1, \dots, L\} \rightarrow \{1, \dots, L\}$, we let σ^{-1} denote its unique inverse permutation. For any vector $a \in \mathbb{R}^L$, $\sigma(a)$ denotes the permutation of a under σ , i.e., $[\sigma(a)]_l = a_{\sigma_l}$. If $a \in \{1, \dots, L\}^L$ is a permutation, then $\sigma(a)$ is also a permutation and we often write σa instead. Let $\mathbf{l} = (1, \dots, L)$ denote the identity permutation. Then $\sigma \mathbf{l} = \sigma$. See [52] for more details. Finally, denote dm_l^j/dp_l by \dot{m}_l^j .

Theorem 2.8. *Suppose assumptions A1–A3 hold. If, for any vector $\mathbf{j} \in \{1, \dots, J\}^L$ and any permutations $\sigma, \mathbf{k}, \mathbf{n}$ in $\{1, \dots, L\}^L$,*

$$\prod_{l=1}^L \dot{m}_l^{[\mathbf{k}(\mathbf{j})]_l} + \prod_{l=1}^L \dot{m}_l^{[\mathbf{n}(\mathbf{j})]_l} \geq \prod_{l=1}^L \dot{m}_l^{[\sigma(\mathbf{j})]_l} \quad (2.25)$$

then the equilibrium of a regular network is globally unique.

Proof. See Appendix 6.2.

Theorem 2.8 implies that if the (slopes of the) price mapping functions are “similar”, then global uniqueness is guaranteed, as the following corollary shows: If \dot{m}_l^j do not differ much across source types at each link, or they do not differ much along links in every source’s path, the equilibrium is unique.

Corollary 2.9. *Suppose assumptions A1–A3 hold. The equilibrium of a regular network is globally unique if any one of the following conditions holds:*

1. For each $l = 1, \dots, L$, $j = 1, \dots, J$

$$\dot{m}_l^j \in \left[a_l, 2^{\frac{1}{L}} a_l \right] \quad \text{for some } a_l > 0 \quad (2.26)$$

2. For each $j = 1, \dots, J$, $l = 1, \dots, L$

$$\dot{m}_l^j \in \left[a^j, 2^{\frac{1}{L}} a^j \right] \quad \text{for some } a^j > 0 \quad (2.27)$$

Proof. If (2.26) holds, we have for any $j_l, \hat{j}_l, \tilde{j}_l$ in $\{1, \dots, J\}$

$$\prod_{l=1}^L \dot{m}_l^{j_l} + \prod_{l=1}^L \dot{m}_l^{\hat{j}_l} \geq 2 \prod_{l=1}^L a_l = \prod_{l=1}^L 2^{\frac{1}{L}} a_l \geq \prod_{l=1}^L \dot{m}_l^{\tilde{j}_l}$$

which implies the sufficient condition in Theorem 2.8.

For the second assertion, fix any j in $\{1, \dots, L\}^L$ and any permutations $\sigma, \mathbf{k}, \mathbf{n}$ in $\{1, \dots, L\}^L$. If (2.27) holds, we have

$$\begin{aligned} \prod_{l=1}^L \dot{m}_l^{[\mathbf{k}(j)]_l} + \prod_{l=1}^L \dot{m}_l^{[\mathbf{n}(j)]_l} &\geq 2 \prod_{l=1}^L a^{j_l} = \prod_{l=1}^L 2^{\frac{1}{L}} a^{j_l} \\ &\geq \prod_{l=1}^L \dot{m}_l^{[\sigma(j)]_l} \end{aligned}$$

which implies the sufficient condition in Theorem 2.8. \square

Remarks:

1. Asymptotically when $L \rightarrow \infty$, both conditions (2.26) and (2.27) converge to a single point. Condition (2.26) reduces to $\dot{m}_l^j = a_l$ which essentially says that all protocols are the same ($J = 1$). Condition (2.27) reduces to $\dot{m}_l^j = a^j$, which is the linear link independent case that will be discussed in Theorem 2.14.
2. These link-based uniqueness results hold for a network whenever no flow uses more than L links.

2.3.4.2 $L = 3$ and $J = 2$ case

We now focus on the case of $L = 3$, $J = 2$ and provide stronger results than a direct application of Theorem 2.8. This case represents the smallest network that can exhibit non-unique equilibrium points if A1–A3 are satisfied and R is full rank (see Theorem 2.17).

Theorem 2.10. *Suppose assumptions A1–A3 hold for a three-links regular network with*

two protocols. If the following six inequalities hold, the network has a unique equilibrium:

$$\lambda_2 + \lambda_3 \geq \lambda_1, \lambda_1 + \lambda_3 \geq \lambda_2, \lambda_1 + \lambda_2 \geq \lambda_3$$

$$\frac{1}{\lambda_2} + \frac{1}{\lambda_3} \geq \frac{1}{\lambda_1}, \frac{1}{\lambda_1} + \frac{1}{\lambda_3} \geq \frac{1}{\lambda_2}, \frac{1}{\lambda_1} + \frac{1}{\lambda_2} \geq \frac{1}{\lambda_3}$$

where $\lambda_l := \dot{m}_l^1(p)/\dot{m}_l^2(p)$.

Proof. See Appendix 6.3.

A straightforward corollary is the following:

Corollary 2.11. *Suppose assumptions A1–A3 hold. For a three-links regular network with two protocols, if, for all l , $\lambda_l \in [a, 2a]$ for some constant $a > 0$, the network admits a globally unique equilibrium.*

Remark: If $\dot{m}_l^j = k^j$ are link independent, then $\lambda_l = k^1/k^2 \in [a, 2a]$ for any $k^1/2k^2 \leq a \leq k^1/k^2$. Hence global uniqueness is guaranteed, which agrees with Theorem 2.14.

We illustrate in Figures 2.5 and 2.6 the regions of λ_l in Theorem 2.10 and Corollary 2.11. They are both cones. The first one is the projection to $\lambda_1 - \lambda_2$ plane and the second one is the cross-section cut by plane $\lambda_1 + \lambda_2 = 1$.




Figure 2.5: Region of λ_l for global uniqueness: projection to $\lambda_1 - \lambda_2$ plane.




Figure 2.6: Region of λ_l for global uniqueness: cross-section cut by plane $\lambda_1 + \lambda_2 = 1$.

2.3.5 Global uniqueness: Jacobian $J(p)$

In a single-protocol network, for the equilibrium price to be unique, it is sufficient that the routing matrix R has full row rank. Otherwise, only the source rates are unique, not necessarily the link prices. In a multi-protocol network, this is no longer sufficient. We now provide another sufficient condition that plays the same role in a multi-protocol network as the rank condition on R does in a single-protocol network (see also the remark after Theorem 2.13).

Let $f = (f_1, \dots, f_n)$ be a vector of real-valued functions defined on \mathfrak{R}^n . Let $G := \{z \in \mathfrak{R}^n | f(z) = 0\}$ and $\text{co}G$ be its convex hull. Define a set $V(G)$ of vectors as

$$V(G) := \{v | v = \phi - \psi \text{ for } \psi, \phi \in \text{co}G\} \quad (2.28)$$

as a function of the set G .

Lemma 2.12. *If for every $z \in \text{co}G$, the Jacobian matrix $J(z) = \partial f(z)/\partial z$ exists and $v^T J(z)v < 0$ for all $v \in V(G)$, then G contains at most one point.*

Proof. For the sake of contradiction, assume there are two distinct points ϕ and ψ in G such that $f(\phi) = f(\psi) = 0$. Let

$$g(\theta) := \phi + \theta(\psi - \phi) \text{ where } \theta \in [0, 1]$$

Then

$$\frac{df(g(\theta))}{d\theta} = \mathbf{J}(g(\theta)) \frac{dg(\theta)}{d\theta} = \mathbf{J}(g(\theta))(\psi - \phi)$$

Hence,

$$f(\psi) - f(\phi) = \int_0^1 \mathbf{J}(g(\theta))(\psi - \phi) d\theta$$

Multiplying both sides by $(\psi - \phi)^T$ yields

$$\begin{aligned} (\psi - \phi)^T (f(\psi) - f(\phi)) &= \\ \int_0^1 (\psi - \phi)^T \mathbf{J}(g(\theta)) (\psi - \phi) d\theta & \end{aligned}$$

The left-hand side of the above equation is 0, and the right-hand side is negative under the assumption of the theorem. This contradiction proves the theorem. \square

Let $f = y$, and let $G = E$ be the set of network equilibria. Then Lemma 2.12, together with Theorem 2.2, provides a sufficient condition for global uniqueness of network equilibrium.

Theorem 2.13. *Suppose assumptions A1–A3 hold. If for every price vector $p \in coE$, the Jacobian matrix $\mathbf{J}(p)$ defined in (2.7) exists and $v^T \mathbf{J}(p)v < 0$ for all $v \in V(E)$, then there exists a globally unique network equilibrium.*

In the single-protocol case, a similar result has been obtained in [55]. However, for that case, the Jacobian matrix is negative definite when R has full row rank. Then the condition in Theorem 2.13 always holds and the equilibrium is unique. In the multi-protocol case, the Jacobian matrix is in general not symmetric and hence not negative definite. Therefore R having full row rank is no longer sufficient for the condition in the theorem to hold.

Since we do not know the equilibrium set E , the condition in the theorem cannot be directly applied to prove global uniqueness. To use the theorem, however, it is sufficient to find a convex superset \tilde{E} of E and a superset \tilde{V} of $V(E)$ such that $v^T \mathbf{J}(p)v < 0$ for all $p \in \tilde{E}$ and $v \in \tilde{V}$. This implies the condition in Theorem 2.13 and hence global

uniqueness. We illustrate this procedure in the next example.

Example 2.4: application of Theorem 2.13 to verify global uniqueness

We visit Example 2.1 for the third time but using log utility functions for all sources, i.e.,

$$U_i^j(x_i^j) = \log(x_i^j) \text{ for all } (j, i) \quad (2.29)$$

Let the Jacobian matrix be

$$\mathbf{J}(p) = \begin{bmatrix} J_{11} & J_{12} & J_{13} \\ J_{21} & J_{22} & J_{23} \\ J_{31} & J_{32} & J_{33} \end{bmatrix}$$

where $J_{kl} = J_{kl}(p)$ are functions of prices p given by (2.7). For example

$$J_{11} = -\frac{1}{p_1^2} - \frac{5}{(5p_1 + p_2)^2} - \frac{1}{(p_1 + p_3 + 3p_2)^2}$$

It can be seen that $\mathbf{J}(p)$ is not negative definite for general p unlike in the single-protocol case. Even though E can be hard to find, we demonstrate how to find a simple convex superset \tilde{E} of E and a simple superset \tilde{V} of $V(E)$.

Consider the convex set

$$\tilde{E} := \{p \in \mathfrak{R}_+^3 \mid 1 \leq p_1 = p_3 \leq 2, 1 \leq p_2 \leq 2\}$$

We claim that $E \subseteq \tilde{E}$. To see this, let p be an equilibrium price. If $p_1 < 1$, then $x_1^1 = 1/p_1$ will exceed the link capacity 1, and hence $p_1 \geq 1$. A similar argument gives $p_2 \geq 1$. To see $p_1 \leq 2$, assume it is not true. Then

$$\begin{aligned} x_1^1 &= 1/p_1 < 1/2 \\ x_1^2 &= 1/(5p_1 + p_2) < 1/11 \\ x_1^3 &= 1/(2p_1 + 3p_2) < 1/7 \end{aligned}$$

Summing them yields $x_1^1 + x_1^2 + x_1^3 < 1$. Hence the network is not in equilibrium, contradicting that p is an equilibrium price. Hence $p_1 \leq 2$. The argument for $p_2 \leq 2$ is similar.

Using the definition of \tilde{E} , we can bound all $J_{kl}(p)$ for $p \in \tilde{E}$. The results are collected in Table 2.3.

Table 2.3: Example 2.4: bounds on elements of $\mathbf{J}(p)$

| Elements | Upperbound | Lowerbound |
|----------|------------|------------|
| J_{11} | -0.2947 | -1.1789 |
| J_{22} | -0.2939 | -1.1756 |
| J_{33} | -0.2947 | -1.1789 |
| J_{23} | -0.0447 | -0.1789 |
| J_{32} | -0.0369 | -0.1478 |
| J_{12} | -0.0369 | -0.1478 |
| J_{21} | -0.0447 | -0.1789 |
| J_{13} | -0.0100 | -0.0400 |
| J_{31} | -0.0100 | -0.0400 |

Let

$$\tilde{V} := \{v \in \mathfrak{R}_+^3 \mid v_1 = v_3\}$$

We claim that $V(E) \subseteq \tilde{V}$. To show this, note that $\text{co}E \subseteq \tilde{E}$ since $\text{co}E$ is the smallest convex set that contains E . Hence $V(E) \subseteq V(\tilde{E})$. Since $p_1 = p_3$ at equilibrium, $v_1 = v_3$ holds for any $v \in V(\tilde{E})$ from the definition of \tilde{E} . Hence, $V(\tilde{E}) \subseteq \tilde{V}$ and therefore $V(E) \subseteq \tilde{V}$.

We now check that $v^T \mathbf{J}(p)v < 0$ for all $p \in \tilde{E}$ and $v \in \tilde{V}$. For any $v \in \tilde{V}$, $v^T \mathbf{J}(p)v$ is the following quadratic form in v_1 and v_2 :

$$\begin{aligned} v^T \mathbf{J}(p)v &= v_1^2(J_{11} + J_{33} + J_{13} + J_{31}) + \\ &\quad v_1 v_2(J_{12} + J_{21} + J_{23} + J_{32}) + v_2^2 J_{22} \end{aligned}$$

If v_1 and v_2 have the same signs, then since J_{kl} are all negative from Table 2.3, $v^T \mathbf{J}(p)v <$

0. If v_1 and v_2 have opposite sign, then a sufficient condition for $v^T \mathbf{J}(p)v < 0$ is

$$(J_{12} + J_{21} + J_{23} + J_{32})^2 < 4J_{22}(J_{11} + J_{33} + J_{13} + J_{31})$$

Using Table 2.3, it is easy to check that the maximum value of $(J_{12} + J_{21} + J_{23} + J_{32})^2 - 4J_{22}(J_{11} + J_{33} + J_{13} + J_{31})$ is -0.2895 . Therefore we have found a superset \tilde{E} of $\text{co}E$ and a superset \tilde{V} of $V(E)$ such that $v^T \mathbf{J}(p)v < 0$ for all $p \in \tilde{E}$ and all $v \in \tilde{V}$. This implies the condition of Theorem 2.13 and hence the global uniqueness of network equilibrium. \square

2.3.6 Global uniqueness of special networks

In this section, we present special networks that have globally unique equilibrium.

2.3.6.1 Case 1: linear link-independent m^j

When the price mapping functions are linear and link-independent, i.e., $m_l^j(p_l) = k^j p_l$ for some scalars $k^j > 0$, it is easy to show that we have an unusual situation in the theory of heterogeneous protocols where the equilibrium rate vector x solves the following concave maximization problem:

$$\begin{aligned} \max_{x \geq 0} \quad & \sum_{i,j} k^j U_i^j(x_i^j) \\ \text{subject to} \quad & Rx \leq c \end{aligned}$$

Therefore, such a network always has a globally unique equilibrium when U_i^j are strictly concave. Here we provide another proof using Theorem 2.13.

Theorem 2.14. *Suppose assumptions A1–A3 hold and R has full row rank. If for all j and l , $m_l^j(p_l) = k^j p_l$ for some scalars $k^j > 0$, then there is a unique network equilibrium.*

Proof. We prove this by showing that the Jacobian matrix $\mathbf{J}(p)$ defined in (2.7) is negative definite over all $p \geq 0$. Then the result follows from Theorem 2.13.

Under the assumptions of the theorem, $\mathbf{J}(p)$ can be simplified into (from (2.7)–(2.9))

$$\begin{aligned}\mathbf{J}(p) &= \sum_j R^j D^j(p) (R^j)^T \frac{\partial m^j(p)}{\partial p} \\ &= \sum_j k^j R^j D^j(p) (R^j)^T\end{aligned}$$

where $D^j(p) = \partial x^j(p)/\partial q^j$. Since U_i^j are strictly concave, $D^j(p)$ is a strictly negative diagonal matrix for all $p \geq 0$. Now, $J(p)$ is symmetric. Moreover, since R has full row rank, RR^T is positive definite, i.e., for any nonzero vector $v \in \mathfrak{R}^L$,

$$\sum_j v^T R^j (R^j)^T v = \sum_j ((R^j)^T v)^T (R^j)^T v > 0$$

Then there exists at least one j such that $\eta^j := (R^j)^T v$ is nonzero. Without loss of generality, assume it is $j = 1$. Then

$$\begin{aligned}v^T \mathbf{J}(p) v &= v^T \sum_j k^j R^j D^j(p) (R^j)^T v \\ &= \sum_j k^j (\eta^j)^T D^j(p) \eta^j \\ &\leq k^1 (\eta^1)^T D^1(p) \eta^1 < 0\end{aligned}$$

where the first inequality follows from the fact that $D^j(p)$ is negative definite. Hence $\mathbf{J}(p)$ is negative definite. \square

2.3.6.2 Case 2: linear network

Consider the classic linear network shown in Figure 2.7.

Theorem 2.15. *Suppose assumptions A1–A2 hold. The linear network in Figure 2.7 has a unique equilibrium.*

Proof. Take $\Lambda = I$ in the gradient algorithm (2.23). We will prove that all the eigenvalues




Figure 2.7: Theorem 2.15: linear network.

of the Jacobian matrix

$$\mathbf{J}(p) = \sum_j R^j D^j(p) (R^j)^T \frac{\partial m^j(p)}{\partial p}$$

have negative real part for all $p \geq 0$. This implies that all equilibria are locally stable. By Corollary 2.7 there must be a unique equilibrium.

In the network shown in Figure 2.7, for $j = 1 \dots L$,

$$(R^j)^T \frac{\partial m^j(p)}{\partial p} = \frac{\partial m_j^j(p)}{\partial p_j} (e^j)^T$$

Since $D^j(p)$ is a negative scalar, we can define a positive number β_j such that:

$$R^j D^j(p) (R^j)^T \frac{\partial m^j(p)}{\partial p} = -\beta_j e^j \cdot (e^j)^T$$

For $j = L + 1$, $\partial m^j(p)/\partial p$ is a positive definite diagonal matrix. Recall that $D^j(p)$ is a scalar. Assume that the i th diagonal entry of matrix $D^j(p) \partial m^j(p)/\partial p$ is $-\gamma_i$. Denote by γ the $L \times 1$ vectors formed from γ_i . Then for $j = L + 1$:

$$R^j D^j(p) (R^j)^T \frac{\partial m^j(p)}{\partial p} = -\mathbf{1} \cdot \mathbf{1}^T \text{diag}(\gamma_i) = -\mathbf{1} \gamma^T$$

By combining the results above, we obtain

$$\begin{aligned}
\mathbf{J}(p) &= \sum_{j=1}^{L+1} R^j D^j (R^j)^T \frac{\partial m^j}{\partial p} \\
&= - \sum_1^L \beta_j e^j \cdot (e^j)^T - \mathbf{1} \gamma^T \\
&= -\text{diag}(\beta_i) - \mathbf{1} \gamma^T
\end{aligned}$$

By the following lemma, all the eigenvalues of above matrix have negative real parts. Therefore, there must be a unique equilibrium by Corollary 2.7. \square

Lemma 2.16. *Suppose that B is a positive definite diagonal matrix, and γ is a positive vector, then the eigenvalues of $B + \mathbf{1} \gamma^T$ have positive real parts.*

Proof: See Appendix 6.5. \square

The theorem can be generalized to include more than one multi-hop flows, provided they all belong to the same type $L + 1$ and the sets of links they traverse are nested, i.e., $L(x_1^{L+1}) \supseteq L(x_2^{L+1}) \supseteq \dots \supseteq L(x_n^{L+1})$ for n multi-hop flows. This result implies that the two 2-link flows in Example 2.1 are necessary to demonstrate non-uniqueness.

Experiment 2.1: linear network: unique equilibrium

We further provide a Dummynet experiment to verify Theorem 2.15 for a three-link network with topology shown in Figure 2.8. The topology is the same as the one used in Experiment 1.1. Each Dummynet router is configured to have 40 ms one-way propagation delay and 200-packet buffer. The link bandwidth is 100 Mbps for link 1, 150 Mbps for link 2, and 120 Mbps for link 3. There are three FAST TCP flows using the paths 1, 2, and 3 with one flow on each path. There are eight Reno flows using path 4. Thirty experiments are done for each scenario.

The average aggregate flow rates and their standard deviations on each of paths 1, 2, 3, and 4 are shown in Table 2.4. They suggest that the network has reached the same equilibrium regardless of which flows start first. This is further confirmed by the queue and throughput trajectories at links 1–3 in Figures 2.9–2.11. At each link, the queue and




Figure 2.8: Experiment 2.1: unique equilibrium.

| | Path 1 (FAST) | Path 2 (FAST) | Path 3 (FAST) | Path 4 (Reno) |
|------------------|------------------|------------------|------------------|------------------|
| FAST start first | (47.8, 2.7) Mbps | (96.2, 2.8) Mbps | (67.2, 2.8) Mbps | (47.9, 2.7) Mbps |
| Reno start first | (46.1, 0.8) Mbps | (94.2, 0.8) Mbps | (64.6, 3.7) Mbps | (43.7, 1.9) Mbps |

Table 2.4: Average aggregate rates and their standard deviations of all flows on paths 1, 2, 3, 4.

throughput behaviors are very similar regardless of whether FAST or Reno flows start first.




Figure 2.9: Experiment 2.1: queue size and aggregate throughput at link 1.

2.3.6.3 Case 3: networks with no flow using more than two links

Theorem 2.7 implies the global uniqueness of equilibrium for any network with no more than 2 links. In this case, the Jacobian matrix $J(p)$ is strictly diagonally dominant with negative diagonal entries, and hence its determinant is $(-1)^L$.




Figure 2.10: Experiment 2.1: queue size and aggregate throughput at link 2.




Figure 2.11: Experiment 2.1: queue size and aggregate throughput at link 3.

Theorem 2.17. *Suppose assumptions A1–A2 hold and R has full row rank. A network that has multiple equilibria must have at least three links.*

Proof. When there is only one link, the Jacobian matrix $\mathbf{J}(p)$ reduces to a negative real number, since $R^j D^j(p)(R^j)^T$ is negative, and $\partial m^j / \partial p$ is positive. Therefore, any equilibrium is locally stable. Hence it is unique by Theorem 2.7.

When there are two links, the Jacobian matrix at any p is

$$\mathbf{J}(p) = \sum_j R^j D^j(p)(R^j)^T \frac{\partial m^j(p)}{\partial p}$$

It can be checked that $\mathbf{J}^T(p)$ is diagonally dominant with strictly negative diagonal entries. Moreover, the full rank condition on R implies that there are sources (j, i) such that $R_{1i}^j R_{2i}^j = 0$, and hence $\mathbf{J}^T(p)$ is *strictly* diagonally dominant. This implies that $\mathbf{J}^T(p)$ is negative definite with strictly negative real eigenvalues. Since $\mathbf{J}(p)$ and $\mathbf{J}^T(p)$ have the same eigenvalues, $\mathbf{J}(p)$ and hence all equilibria are locally stable. According to Theorem 2.7, there is a unique equilibrium. \square

Remark: If R does not have full row rank, then there are two-link networks that have multiple equilibria; see Appendix 6.4.

2.4 Related Work in Economics

Our formulation is close to the general equilibrium theory in economics from which we borrow ideas and techniques [51]. See [18, 21, 22, 57, 58, 78, 17, 19] and [49, 6] for a fairly complete treatment of related works in economics literature. A typical model of the pure exchange economy consists of L commodities and N consumers. Each consumer i has an initial endowment vector $\omega_i = (\omega_{il} \geq 0, l = 1, \dots, L)$ and its goal is to choose a consumption vector $x_i = (x_{il}, l = 1, \dots, L)$ to maximize its utility subject to its wealth constraint, i.e.,

$$\max_{x_i \geq 0} U_i(x_i) \quad \text{subject to} \quad p^T x_i \leq p^T \omega_i$$

where $p = (p_l, l = 1, \dots, L)$ are unit prices for the goods and T denotes matrix transpose. For each good $l = 1, \dots, L$, demand and supply are balanced if

$$\sum_{i=1}^N x_{il} = \sum_{i=1}^N \omega_{il}$$

A consumption vector $x^* = (x_i^*, i = 1, \dots, N)$ and a price vector p^* are called a *competitive equilibrium* (or *Walrasian equilibrium*) if x_i^* maximizes i 's utility and demand equals supply for all goods.

In general equilibrium theory, consumers are assumed to be price takers. This aspect is similar to our model where sources do not take into account how their decisions affect the link prices or each other. Both problems are concerned with characterizing fixed points of a continuous mapping, and hence there are considerably similarities in terms of the characterizations and the mathematical tools to derive them. The main mathematical tools used in this paper are the Nash theorem in game theory [56, 9], which is an application of Kakutani's generalized fixed point theorem, and results from differential topology, especially the Poincare-Hopf index theorem [53]. They are used to prove existence and study uniqueness of network equilibrium, respectively. There are however several differences.

First, the effective prices to different sources (consumers) are generally different in our model, whereas the prices in the economic model are independent of consumers. Differential pricing is what makes networks with heterogeneous protocols much more difficult. Second, in the economic model, there is a concept of initial endowment that defines both the demand-supply relation and a consumer's consumption possibility through the wealth constraint. In our model, the wealth constraint is replaced by the link capacity constraint. Third, in the economic model, consumers maximize their utilities whereas in our model, sources maximize their utilities minus bandwidth costs. Finally, in our model, every source consumes exactly the same amount of bandwidth at each link in its path ($x_{il} = x_i$, for all $l \in L(i)$), whereas, in the economic model, consumers can consume different goods at different amounts. This guarantees that the demand for every good is exactly balanced by its supply in a pure exchange economy; yet in networks, the set of bottleneck links where demand for and supply of bandwidth is balanced can be non-unique and a strict subset of all

links. The property $x_{il} = x_i$ is the key structure that allows us to obtain interesting results on global uniqueness in fairly general settings. In contrast, global uniqueness in general equilibrium analysis usually requires very strong conditions and most literature focuses on local uniqueness [17, 19, 6]. We will return to related problems in general equilibrium theory at the end of chapter 5.

Chapter 3

Optimality and Stability

A central issue in networking is how to allocate bandwidth to flows *efficiently* and *fairly*, in a decentralized manner. A series of recent work, e.g. [39, 46, 55, 84, 47, 41, 44], has shown that a bandwidth allocation policy can be expressed in terms of a utility function $U_i(x_i)$ in the sense that the desired bandwidth allocation $x^* = (x_i^*, \text{all sources } i)$ solves the utility maximization problem (2.10)–(2.11). As shown in section 2.1.3, (2.10)–(2.11) characterize equilibrium completely for homogeneous price case and the optimality here includes guarantee for both efficiency and fairness. However, as we have shown in chapter 2, for heterogeneous congestion control networks, equilibrium can not be characterized by (2.10)–(2.11) anymore. In this chapter, we first look at the deviation of efficiency and fairness in section 3.1. In terms of efficiency, it is shown that qualitatively equilibrium is still Pareto efficient but quantitatively there is efficiency loss about which we can provide an upperbound. On fairness, we show that intra-protocol fairness is still decided by utility maximization problem while inter-protocol fairness is the part which we don't have control over. However it is shown that we can achieve any desirable inter-protocol fairness by properly choosing protocol parameters. This analysis provides insights on networks with heterogeneous congestion signals and further motivates the algorithm design in chapter 4. For an engineering system, stability of equilibrium is essential. We investigate it in section 3.2 and in particular show that if the degree of pricing heterogeneity is small enough, then not only is the network equilibrium unique, it is also locally stable.

3.1 Optimality

3.1.1 Efficiency

In this section, efficiency of equilibrium of networks with heterogeneous protocols is explored. We first make the following key observation, which not only leads to the remaining results of this subsection, but also will be the starting point of our algorithm design in Chapter 4.

3.1.1.1 Qualitative results: Pareto efficiency

Theorem 3.1. *Given an equilibrium p^* , there exists a positive vector γ , such that the equilibrium rate vector $x^*(p)$ is the unique solution of following problem:*

$$\max_{x \geq 0} \sum_{i,j} \gamma_i^j U_i^j(x_i^j) \quad (3.1)$$

$$\text{subject to} \quad Rx \leq c \quad (3.2)$$

Proof. The KKT (Karush-Kuhn-Tucker) optimality conditions for (3.1), (3.2) are:

$$\gamma_i^j (U_i^j)'(x_i^j) = \sum_l R_{il}^j p_l \quad \text{for all } (i, j) \quad (3.3)$$

$$p^T (Rx - c) = 0 \quad (3.4)$$

$$Rx - c \leq 0 \quad (3.5)$$

where the (x, p) are the primal-dual variables. We now claim our system satisfies those conditions with equilibrium rates and prices (x^*, p^*) by choosing

$$\gamma_i^j = \frac{\sum_l R_{il}^j p_l^*}{\sum_l R_{il}^j m_l^j(p_l^*)} \quad (3.6)$$

To see this, note (3.4) and (3.5) are conditions for equilibrium. After substituting (3.6) into

(3.3), we have

$$(U_i^j)'(x_i^{j*}) = \sum_l R_{il}^j m_l^j(p_i^*) \quad (3.7)$$

That is also an equation used to define equilibrium. \square

It is worthwhile to note that Theorem 3.1 gives an underlying convex optimization problem an equilibrium solves, but this optimization problem itself depends on equilibrium. Hence it cannot be used to find equilibrium directly, nor does it give existence and uniqueness as in the single-protocol case [76].

As stated by the celebrated first fundamental theorem of welfare economics, any competitive equilibrium is Pareto efficient. That explains the most basic reason that congestion signals are used to regulate source rates and hence realize bandwidth allocation. We know the unique equilibrium is Pareto efficient when there is a single price. Now we can show that the same holds for networks with heterogeneous protocols as a direct corollary of Theorem 3.1:

Corollary 3.2. *All equilibrium points are Pareto efficient.*

3.1.1.2 Quantitative results: price of heterogeneity

Pareto efficiency can be viewed as a qualitative requirement for an optimal allocation. However, it does not give a quantitative criterion for optimum. Aggregate utility (social welfare) is the standard criterion for optimum. As shown in Section 2.1.3, when there is only one price, the unique equilibrium achieves the maximum aggregate utility. For heterogeneous protocols cases, we now study efficiency loss by lower-bounding the ratio of the achieved aggregate utility to its maximum.

Theorem 3.3. *Assume all utility functions are nonnegative, i.e., $U(x) \geq 0$. Suppose the optimal aggregate utility is U^* and \hat{U} is the achieved aggregate utility at equilibrium of a network with heterogeneous protocols. Then*

$$\frac{\hat{U}}{U^*} \geq \frac{\gamma}{\bar{\gamma}} \quad (3.8)$$

where $\underline{\gamma}$ and $\bar{\gamma}$ are the lower and upper bounds of γ_i^j ¹.

Proof. Assume \hat{x} is one of the solutions of Theorem 3.1, then

$$\max_x \sum_{i,j} \gamma_i^j U_i^j(x_i^j) = \sum_{i,j} \gamma_i^j U_i^j(\hat{x}_i^j) \leq \bar{\gamma} \hat{U} \quad (3.9)$$

On the other hand,

$$\max_x \sum_{i,j} \gamma_i^j U_i^j(x_i^j) \geq \underline{\gamma} \max_x \sum_{i,j} U_i^j(x_i^j) = \underline{\gamma} U^* \quad (3.10)$$

Combining the two equalities above, we get

$$\frac{\hat{U}}{U^*} \geq \frac{\underline{\gamma}}{\bar{\gamma}}$$

□

It has been known for a long time that price can serve as the “invisible hand” to coordinate competing users and realize optimal resource allocation. That however requires two basic assumptions. The first assumption is that users are all price takers. If instead they are noncooperative game players, there will be efficiency loss. Such “price of anarchy” was recently bounded from above for both routing [62] and congestion control [36]. The second assumption is the homogeneity of price that all users see, which does not hold in networks with more than one type of congestion control protocols. Our result above quantifies the efficiency loss as a “price of heterogeneity”.

The following simple example is used here to show the efficiency loss can be arbitrarily large. Consider two flows sharing a single link with unit capacity. Flow 1’s utility is $\beta_1 \sqrt{x_1}$ and flow 2’s is $\beta_2 \sqrt{x_2}$ where $\beta_1 > \beta_2 > 0$. Suppose flow 2 reacts to price p^2 while flow 1 reacts to $p^1 = kp^2$, in other words, $m^1(p) = kp$ and $m^2(p) = p$. It is straightforward to

¹Both $\underline{\gamma}$ and $\bar{\gamma}$ can be bounded using m_i^j . For example, for a network with both loss based and delay based protocols and assuming RED is used, the slopes of RED at different links can provide information on m_i^j .

calculate that

$$U^* = \sqrt{(\beta_1)^2 + (\beta_2)^2}$$

and

$$\hat{U} = \frac{(\beta_1)^2 + k(\beta_1)^2}{\sqrt{(\beta_1)^2 + k^2(\beta_2)^2}}$$

When $k \rightarrow \infty$, $\hat{U} \rightarrow \beta_2$ and

$$\frac{\hat{U}}{U^*} \rightarrow \frac{1}{1 + (\frac{\beta_1}{\beta_2})^2}$$

Therefore $\frac{\hat{U}}{U^*}$ can be arbitrarily close to zero when $\frac{\beta_1}{\beta_2}$ is sufficiently large.

3.1.2 Fairness

In this section, we study fairness in networks shared by heterogeneous congestion control protocols. Two questions we address are: how the flows within each protocol share among themselves (intra-protocol fairness) and how these protocols share bandwidth in equilibrium (inter-protocol fairness).

3.1.2.1 Intra-protocol fairness

As indicated above, when the network is shared only by flows using the same protocol, the equilibrium flow rates are the unique optimal solution of a utility maximization problem. In other words, the utility functions describe how the flows share bandwidth among themselves. For instance, the log utility function of FAST implies that it achieves weighted proportional fairness. When flows using different congestion signals share the same network, it turns out that this feature is still preserved “locally” within each protocol, as we now show. In particular, e.g., it implies that the intra-protocol fairness of FAST is still proportional fairness.

Theorem 3.4. *Given an equilibrium $(\hat{x}, \hat{p}) \geq 0$, let $\hat{c}^j := R^j \hat{x}^j$ be the total bandwidth consumed by flows using protocol j at each link. The corresponding flow rates \hat{x}^j are the unique solution of:*

$$\max_{x^j \geq 0} \sum_{i=1}^{N^j} U_i^j(x_i^j) \quad \text{subject to } R^j x^j \leq \hat{c}^j \quad (3.11)$$

Proof: Since $(\hat{x}^j, \hat{p}^j) \geq 0$ is an equilibrium, from (2.2) to (2.5), we have

$$(U_i^j)'(\hat{x}_i^j) = \sum_l R_{li}^j \hat{p}_l^j \quad \text{for } i = 1, \dots, N^j$$

This, together with (from the definition of \hat{c}^j)

$$\sum_i R_{li}^j \hat{x}_i^j \leq \hat{c}_l^j, \quad \hat{p}_l^j \left(\sum_i R_{li}^j \hat{x}_i^j - \hat{c}_l^j \right) = 0, \quad \forall l$$

forms the necessary and sufficient condition for \hat{x}^j and \hat{p}^j to be optimal for (3.11) and its dual respectively. \square

Note that in Theorem 3.4, the “effective capacities” \hat{c}^j ’s are not preassigned. They are the outcome of competition among flows using different congestion prices and are related to inter-protocol fairness, which we now discuss.

3.1.2.2 Inter-protocol fairness

Even though flows using different congestion signals individually solve a utility maximization problem to determine their intra-protocol fairness, they in general do not jointly solve any convex utility maximization problem. This makes the study of inter-protocol fairness hard. Here we provide a feasibility result, which says any reasonable inter-protocol fairness is achievable by linearly scaling congestion control algorithms.

Assume flow (j, i) has a parameter μ_i^j with which it decides its rate in the following

way:

$$x_i^j(q_i^j) = (U_i^j)^{\prime-1} \left(\frac{1}{\mu_i^j} q_i^j \right)$$

For example, if we consider FAST's utility function $\alpha \log(x)$, then the α parameter in the protocol can be viewed as μ here. Our main result in this subsection says for a network with J protocols, if $J - 1$ protocols have their linear scalar vectors μ^j , then there exists a μ vector such that one of the resulting equilibria with that μ can achieve any predefined bandwidth partition. Before we get to the theorem itself, we first characterize the feasible set of predefined bandwidth allocation.

Assume that except for $j = J$, flow (i, j) has parameter μ_i^j . Or equivalently, we can define $\mu_i^J = 1$. The equilibrium rates x^j clearly depend on parameter μ . For $j = 1, 2, \dots, J - 1$, let $\bar{x}^j(\mu)$ be the unique rates of flows using protocol j if there were no other protocols in the network. Let $\underline{x}^j(\mu)$ be the unique rates of type j flows if network capacity were $(c - \sum_{k \neq j} R^k \bar{x}^k)^+$.

Let

$$X^* := \{ x, | \underline{x}^j(\mu) \leq x^j \leq \bar{x}^j(\mu), \mu \geq 0, Rx \leq c \}$$

X^* includes all possible rates of flows using protocol j if they were given strict priority over other flows or if others were given strict priority over them, and all rates in between. In this sense X^* contains the entire spectrum of inter-protocol fairness among different protocols. The next result says that every point in this spectrum is achievable by an appropriate choice of parameter μ .

Let $x^j(\mu)$ denote the equilibrium rates (may not be unique) of flows using price j sharing the same network (R, c) with other protocols when the protocol parameter is μ .

Theorem 3.5. *For every link l , assume there is at least one type J flow that only uses that link. Given any $x^* \in X^*$, there exists an $\mu^* \geq 0$ such that $x^j(\mu^*) = (x^j)^*$ for $j = 1, 2, \dots, J - 1$.*

Proof: Given any $x^* \in X^*$, the capacity for all type J flows is $c - \sum_{k \neq J} R^k (x^k)^*$. Since

$Rx^* \leq c$ (for all coordinates), we have $c - \sum_{k \neq J} R^k (x^k)^* \geq (x^J)^*$, which is greater than or equal to 0. Hence the following utility maximization problem solved by flows of type J is feasible:

$$\begin{aligned} & \max_{x^J \geq 0} \quad \sum_i U_i^J(x_i^J) \\ \text{subject to} \quad & R^J x^J \leq c - \sum_{k \neq J} R^k (x^k)^* \end{aligned}$$

Let p^J be an associated Lagrange multiplier vector. By the assumption that every link has at least one single-link type J flow, we know $p_l^J > 0$ for all l . Choose $(\mu^j)^*$ with $(\mu^j)_i^* = (x^j)_i^* \sum_l R_{li}^j m_l^j ((m^J)_l^{-1}(p_l^J))$. It can be checked that all equilibrium equations are satisfied. \square

Remarks:

1. In general, one can view Theorem 3.1 as defining fairness of flows using heterogeneous protocols and can conclude that price mapping functions (router parameters) affect fairness. This is very different from the case when there is only one type of protocol and we will return to this point in Chapter 4. Clearly, if one can choose price mapping functions, one can achieve any predefined fairness. More interesting, Theorem 3.5 shows that we can achieve any reasonable predefined fairness without modifying router parameters but only by choosing a linear scalar in source algorithm. This opens up the possibility of maintaining the end-to-end feature of TCP.
2. Theorem 3.5 implies that given any reasonable fairness among flows using different congestion signals, in terms of a desirable rate allocation x^* , there exists a protocol parameter vector μ^* that achieves it. It is however yet unclear how to compute μ^* dynamically in practice using only local information. In chapter 4, we will discuss distributed algorithms to compute a particular μ^* , which will result in optimal bandwidth allocation.

3.2 Stability

3.2.1 Low dimensional results

For general dynamical systems, a globally unique equilibrium point may not even be locally stable [65, 38]. We start with some low dimensional cases. For networks with heterogeneous protocols, we prove that global uniqueness implies local stability for networks with no more than three bottleneck links, we also prove global stability for networks with no more than two bottleneck links. By the end of this section, We proceed to general networks with any number of links.

Theorem 3.6. *For a network with $L \leq 3$, if there is only one equilibrium, it is also locally stable.*

Proof. We want to prove all eigenvalues of $\mathbf{J}(p)$ lie in the left half plane, where $\mathbf{J}(p)$ is the Jacobian of equilibrium equations ($\mathbf{J}(p) = \frac{\partial \mathbf{y}}{\partial \mathbf{p}}$) evaluated at equilibrium. By index theorem and global uniqueness analysis [76], we have $\det(-\mathbf{J}) > 0$ for the unique equilibrium.

When L equals 1 or 2, it is obvious as $J_{ii} < J_{ij} < 0$ for $j \neq i$. Let's consider the case with $L = 3$. Suppose $\lambda^3 + \rho_1 \lambda^2 + \rho_2 \lambda + \rho_3 = 0$ is the characteristic equation for \mathbf{J} . Then ρ_1 is the trace of $-\mathbf{J}$, ρ_2 is sum of all 2×2 principle minors of $-\mathbf{J}$ and $\rho_3 = \det(-\mathbf{J})$.

The Routh array for the equation is

$$\begin{bmatrix} 1 & \rho_2 \\ \rho_1 & \rho_3 \\ (\rho_1 \rho_2 - \rho_3)/\rho_1 & 0 \\ \rho_3 & 0 \end{bmatrix}$$

Applying Routh stability criterion [26], we need all quantities in the left column to be positive to guarantee all roots lie in the left half plane. Clearly $\rho_1 > 0$. Global uniqueness implies $\rho_3 > 0$. Hence we only need to check $\rho_1 \rho_2 > \rho_3$. This is true because

$$\begin{aligned}
\det(\mathbf{J}) &= J_{11}(J_{22}J_{33} - J_{23}J_{32}) - J_{12}(J_{21}J_{33} - J_{23}J_{31}) \\
&\quad + J_{13}(J_{21}J_{32} - J_{22}J_{31}) \\
&> J_{11}J_{22}J_{33} - J_{11}J_{23}J_{32} + J_{11}J_{22}J_{33} \\
&\quad - J_{12}J_{21}J_{33} + J_{11}J_{22}J_{33} - J_{13}J_{22}J_{31} \\
&= J_{11}(J_{22}J_{33} - J_{23}J_{32}) + J_{22}(J_{11}J_{33} - J_{13}J_{31}) \\
&\quad + J_{33}(J_{11}J_{22} - J_{12}J_{21}) \\
&> (J_{11} + J_{22} + J_{33})(J_{11}J_{22} - J_{12}J_{21}) \\
&\quad + (J_{11}J_{33} - J_{13}J_{31}) + (J_{22}J_{33} - J_{23}J_{32}) \\
&= -\rho_1\rho_2
\end{aligned}$$

The first inequality follows from $J_{ii} < J_{ij} < 0$ for $j \neq i$. The second one follows from $J_{ii}J_{jj} - J_{ij}J_{ji} > 0$ for $j \neq i$. One is referred to [76] for detail properties of \mathbf{J} .

Therefore

$$\rho_3 = \det(-\mathbf{J}) = -\det(\mathbf{J}) < \rho_1\rho_2$$

□

As reviewed in 2.1.3, when there is only one kind of price, global stability is proved by using the objective function of the dual of the system problem as a Lyapunov function. For heterogeneous protocols, we have the following.

Theorem 3.7. *For a network with $L \leq 2$, the equilibrium is globally asymptotically stable.*

Proof. The uniqueness of the equilibrium for a network with no more than two bottleneck links is shown in section 2.3.6.3. Assume the equilibrium price is p^* . We now prove global stability.

When $L=1$, consider the simple quadratic Lyapunov function

$$L(p(t)) = (p(t) - p^*)^2$$

Clearly, $L(p(t)) \geq 0$ and $L(p^*) = 0$. We now evaluate its derivative

$$\dot{L} = 2(p - p^*)\dot{p} = 2\gamma(p - p^*)(y(p) - c_l) < 0$$

Hence the equilibrium is globally asymptotically stable.

When $L=2$, let

$$L(p(t)) = |p_1(t) - p_1^*| + |p_2(t) - p_2^*|$$

We have $L(p(t)) \geq 0$ and $L(p^*) = 0$. If $p_1(t) \leq p_1^*$ and $p_2(t) \leq p_2^*$, we have $y_1(t) \geq c_1$ and $y_2(t) \geq c_2$. Therefore $\dot{L} = -(\dot{p}_1(t) + \dot{p}_2(t)) = -(y_1(t) - c_1) + (y_2(t) - c_2) \leq 0$. Similarly, when $p_1(t) \geq p_1^*$ and $p_2(t) \geq p_2^*$, $\dot{L} \leq 0$. Now consider $p_1(t) \leq p_1^*$ and $p_2(t) \geq p_2^*$, $L(p(t)) = p_1^* - p_1(t) + p_2(t) - p_2^*$. It then follows that $\dot{L} = -\dot{p}_1(t) + \dot{p}_2(t) = -(y_1(t) - c_1) + (y_2(t) - c_2)$. Let y_{11}, y_{22} and y_{12} be the aggregate rates of flows using only link 1, only link 2 and both link 1 and 2. Then $y_1 = y_{11} + y_{12}$ and $y_2 = y_{22} + y_{12}$.

$$\begin{aligned} \dot{L} &= -(y_{11} + y_{12} - y_{11}^* - y_{12}^*) + (y_{22} + y_{12} - y_{22}^* - y_{12}^*) \\ &= -(y_{11} - y_{11}^*) + (y_{22} - y_{22}^*) \leq 0 \end{aligned} \tag{3.12}$$

□

3.2.2 General case

We now state the general result on local stability in the following theorem. It essentially says that if the similarity condition on price mapping functions that guarantees uniqueness (theorem 2.8) is satisfied, the unique equilibrium is also locally stable.

Theorem 3.8. *Suppose assumptions A1–A2 hold. If, for any vector $\mathbf{j} \in \{1, \dots, J\}^L$ and*

any permutations $\sigma, \mathbf{k}, \mathbf{n}$ in $\{1, \dots, L\}^L$,

$$\prod_{l=1}^L \dot{m}_l^{[\mathbf{k}(j)]_l} + \prod_{l=1}^L \dot{m}_l^{[\mathbf{n}(j)]_l} \geq \prod_{l=1}^L \dot{m}_l^{[\sigma(j)]_l} \quad (3.13)$$

then the equilibrium of a regular network is locally stable.

Proof. See Appendix 6.6.

Theorem 3.9. *For every equilibrium and any of its neighborhood, there is at least one point outside the neighborhood, starting from which the trajectory goes into the neighborhood.*

Proof. Consider the vector field generated by

$$\dot{p} = f(p)$$

For any equilibrium p^* and its neighborhood V with boundary ∂V . Divergence theorem says the following:

$$\int_V (\nabla \cdot f) dV = \int_{\partial V} (f \cdot n) da \quad (3.14)$$

Here, dV is "volume" element and da is "area" element with unit outward normal n . It is straightforward to see that

$$\nabla \cdot f = \sum_{l=1}^L \frac{\partial f_l}{\partial p_l} < 0 \quad (3.15)$$

Hence $f \cdot n$ has to be negative somewhere on ∂V , i.e., at those places the trajectories point inwards.

□

Chapter 4

A Slow Timescale Update

As shown in the analysis in Chapters 2 and 3, because of the lack of a global coordinate signal, networks with heterogeneous congestion control protocols have complicated behaviors (e.g. multiple equilibria; see section 2.3) and may suffer from efficiency loss and unpredictable fairness (section 3.1). On the other hand, Theorem 3.1 and 3.5 suggest that properly setting a linear scaling factor in source algorithms is enough to achieve optimality (best efficiency and fairness). These motivate the work in this chapter in which we try to answer the following question in the affirmative. Can we steer a heterogeneous network to the equilibrium point which is optimal in the sense of maximizing aggregate utility in a distributed and stable manner? We propose a general algorithm that is simple, scalable, and deployable to achieve this goal. The basic idea is simple. Besides regulating their rates according to their congestion signals, sources also adapt a parameter in a *slow timescale* based on a common congestion signal. This allows a source to choose a particular congestion signal in a fast timescale (and therefore maintain benefits associated with it) while asymptotically reach the optimal equilibrium. The roadmap of this chapter is as follows. Two experiments involving Reno and FAST are reported in section 4.1 which imply that we cannot predict, nor control, the bandwidth allocation through just the design of congestion control algorithms. In section 4.2, a simple source-based algorithm is presented that decouples bandwidth allocation from network parameters and flow arrival patterns in a heterogeneous network. The rest of the chapter is then devoted to the general framework of this slow timescale control. Analysis on existence of a unique, optimal and stable equilibrium of this algorithm is provided (section 4.3). Numerical examples are used to demonstrate

its correctness and convergence in different operation environments (section 4.4). Finally, more realistic experiments that are conducted using WAN in Lab are reported to show the algorithm's effectiveness and some of its byproducts (section 4.5).¹

4.1 Two Examples

In this section, we describe two experiments to illustrate some bandwidth allocation problems in heterogenous networks. In the next section, we describe a simple algorithm that solves these problems. Both the problems and the solution will be greatly extended to general networks and protocols in the following sections.

Both experiments use Reno TCP, which uses packet loss as congestion signal, and FAST TCP, which uses queueing delay as congestion signal. The first experiment shows that when a Reno flow shares a single bottleneck link with a FAST flow, the relative bandwidth allocations depend critically on the link parameter (buffer size): The Reno flow achieves much higher bandwidth than FAST when the buffer size is large and much smaller bandwidth when it is small. This implies that one cannot control the fairness between Reno and FAST through just the design of TCP congestion control algorithms, since fairness is now linked to network parameters, unlike the case of homogeneous networks.

The second experiment shows that even on a fixed (multi-link) network, one cannot control the fairness between Reno and FAST because the relative allocation changes depending on which flow start first!

4.1.1 Dependence of bandwidth allocation on buffer size

Experiment 4.1a: dependence of bandwidth allocation on buffer size

In this example, one FAST flow and one Reno flow share a single bottleneck link with capacity of 8.3 pkt per ms (equivalent to 100Mbps with standard packet size) and round trip propagation delay 50ms. The topology is shown in Figure 4.1. The FAST flow fixes its α parameter at 50 packets. This means that the bottleneck bandwidth will be equally

¹Most results in this chapter are based on a working paper [77].

shared between FAST and Reno when the buffer size B is 150 packets.

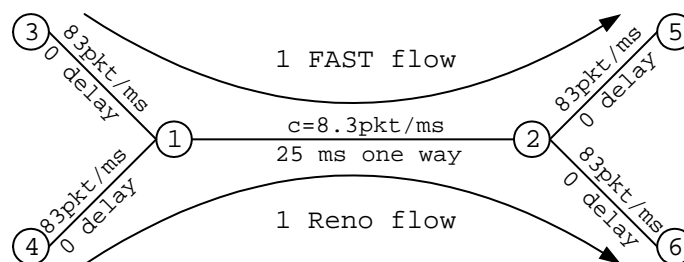


Figure 4.1: Single link example.

In all of the NS-2 simulations in this chapter, heavy-tail noise traffic is introduced in each link with an average rate of 10% of the link capacity.² Figure 4.2 shows the result with a bottleneck buffer size $B = 400$ packets. In this case, FAST gets an average of 2.1 pkt per ms while Reno gets 5.4 pkt per ms. Figure 4.3 shows the result with $B = 80$ packets. Since the bottleneck buffer size is smaller, the average queue in the bottleneck is smaller. Therefore FAST gets a much higher throughput of 3.4 pkt per ms and Reno gets a lower throughput of 0.6 pkt per ms. In this case, the loss rate is fairly high and the aggregate throughput is lower than the bottleneck capacity due to many timeout events.

In summary, bandwidth sharing between Reno and FAST depends on network parameters in a heterogeneous network, contrary to the case of homogeneous network. This is undesirable since the bandwidth allocation among all competing flows in a network should depend only on their valuation of bandwidth (utility functions) but not on AQM parameters. In the next section, we propose a simple source-based solution to achieve this.

4.1.2 Dependence of bandwidth allocation on flow arrival pattern

Experiment 4.2a: dependence of bandwidth allocation on flow arrival pattern

The topology of this network is shown in Figure 4.4. We use RED algorithm [25] and

²The sample figure shows the rate trajectory in one simulation run. The rate value is measured every 2 seconds. The summary figure presents the rate trajectory averaged over 20 simulation runs with different random seeds. Each point in the summary figure represents the average throughput over a period of one minute. The error bars are also shown in the figure.

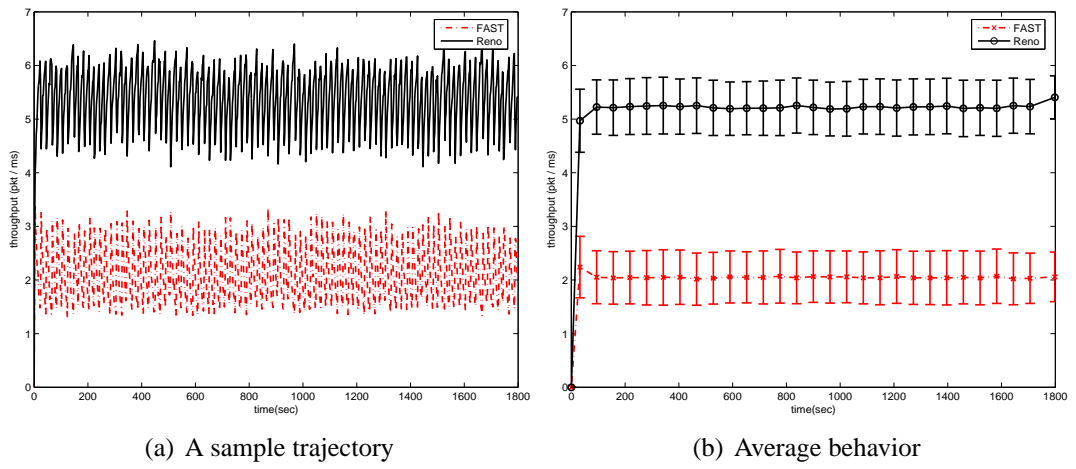


Figure 4.2: FAST vs. Reno with a buffer size of 400 pkts.

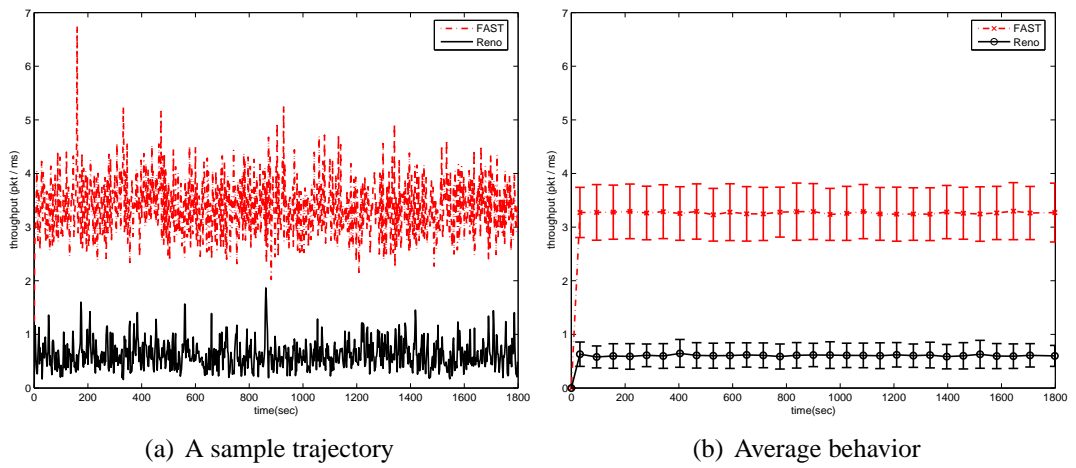


Figure 4.3: FAST vs. Reno with a buffer size of 80 pkts.

packet marking instead of dropping. The marking probability $p(b)$ of RED is a function of queue length b :

$$p(b) = \begin{cases} 0 & b \leq \underline{b} \\ \frac{1}{K} \frac{b-\underline{b}}{\bar{b}-\underline{b}} & \underline{b} \leq b \leq \bar{b} \\ \frac{1}{K} & b \geq \bar{b} \end{cases} \quad (4.1)$$

where \underline{b} , \bar{b} and K are RED parameters. Links 1-2 and 3-4 are both configured with 9.1 pkts per ms capacity (equivalent to 111 Mbps), 30 ms one-way propagation delay, and a buffer of 1500 packets. Their RED parameters are $(\underline{b}, \bar{b}, K) = (300, 1500, 10000)$. Link 2-3 has a capacity of 13.8 pkts per ms (166 Mbps) with 30 ms one-way propagation delay and a buffer size of 1500 packets. Its RED parameters are set to $(0, 1500, 10)$.

There are eight Reno flows on path 1-2-3-4, utilizing all three links, with one-way propagation delay of 90 ms. There are two FAST flows on each of paths 1-2-3 and 2-3-4. Both of them have one-way propagation delay of 60 ms. All FAST flows use a common $\alpha = 50$ packets.




Figure 4.4: Multiple equilibria scenario.

Two sets of simulations have been carried out with different starting times for Reno and FAST flows. The intuition is that if FAST flows start first, link 2-3 will be saturated and links 1-2 and 3-4 will not. Since the RED dropping slope of link 2-3 is steep, when Reno flows join, they will experience so many losses that links 1-2 and 3-4 will remain unsaturated. If Reno flows start first, on the other hand, links 1-2 and 3-4 are saturated while link 2-3 is not because link 2-3 has a higher capacity. Since the RED dropping slopes of link 1-2 and 3-4 are not steep, they can generate enough queueing delay to squeeze FAST flows when they join and keep link 2-3 unsaturated. In the simulations, one set of flows

(Reno or FAST) start at time zero, and the other set of flows start at the 100th second. We present the throughput achieved by one of the FAST flows and one of the Reno flows. Each point in the summary figures represents the average rate over 5 minutes.

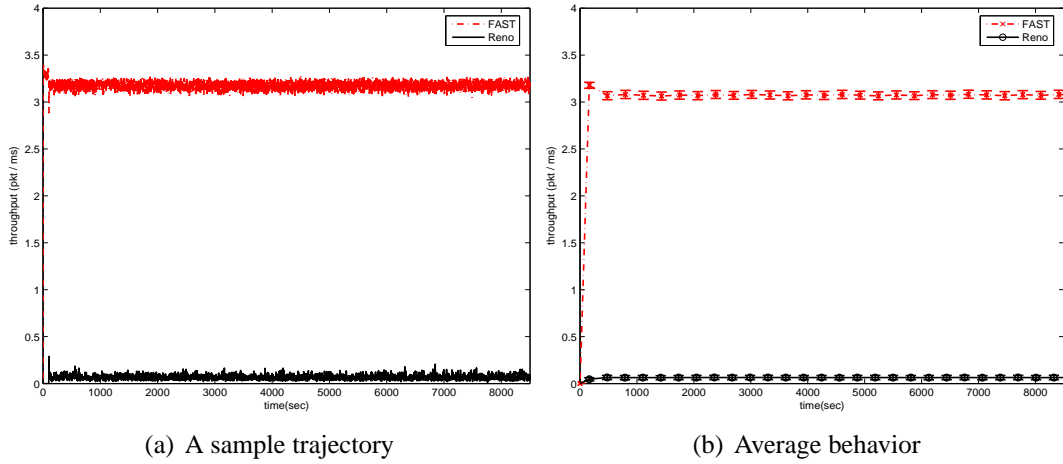


Figure 4.5: Bandwidth shares of Reno and FAST when FAST starts first.




Figure 4.6: Bandwidth shares of Reno and FAST when Reno starts first.

4.2 One Solution

We propose a simple source-based algorithm to solve the problems on unfairness and unpredictable parameter sensitivity illustrated by the examples in the last section. Complete development, theoretical confirmation, and simulation verification of the solution form the rest of the paper after this section. The solution meets all of the following requirements:

Algorithm 1 α adaptation algorithm

1. Every α update interval (2 minutes by default), calculate:

$$\alpha^* = \frac{q}{lw}$$

q and l are average queueing delay and average packet loss rate over the α update interval, w is a parameter. Then

$$\alpha = \begin{cases} \min \{1.1\alpha, \alpha^*\} & \text{if } \alpha < \alpha^* \\ \max \{0.9\alpha, \alpha^*\} & \text{if } \alpha > \alpha^* \end{cases}$$

2. Every window update interval (20ms by default), run FAST algorithm (1.2).
-

1. Always achieve a unique equilibrium that efficiently utilizes the bandwidth.
2. Maintain good fairness between different flows using different protocols.
3. Allocation is independent of AQM setting.
4. Only use end-to-end local information that is available to each flow.
5. Only require simple parameter updates, such as the linear parameter α in FAST.

This α adaptation algorithm, Algorithm 1, fine-tunes the value of α according to the signal of queue and loss in a large time scale (several RTTs). The basic idea of the solution is that FAST should adjust its aggressiveness (α) to the proper level by looking at the ratio of end-to-end queueing delay and end-to-end loss. In other words, FAST also reacts to loss in the slow timescale.

4.2.1 Independence of bandwidth allocation on buffer size

Experiment 4.1b: independence of bandwidth allocation on buffer size

We repeat the simulations in experiment 4.1a with Algorithm 1, w is set to be 125s. Figure 4.7, Figure 4.8, Figure 4.10 and Figure 4.11 should be compared with Figure 4.2, Figure 4.3, Figure 4.5 and Figure 4.6 correspondingly.

With Algorithm 1, FAST achieves 3.4 pkt per ms with buffer size of 400 and 3.2 pkt per ms with buffer size of 80, while Reno gets 4.2 pkt per ms and 4.1 pkt per ms, respectively.

The fairness is greatly improved and only slightly depends on buffer size now, which we summarize in table 4.1 by listing the ratio of bandwidth that Reno gets and FAST gets in different scenarios.

| | B=400 | B=80 |
|---------------------|-------------|--------------|
| Without Algorithm 1 | 5.4/2.1=2.6 | 0.6/3.4=0.18 |
| With Algorithm 1 | 4.2/3.1=1.4 | 4.1/3.2=1.3 |

Table 4.1: Ratio of Reno's rate and FAST's rate

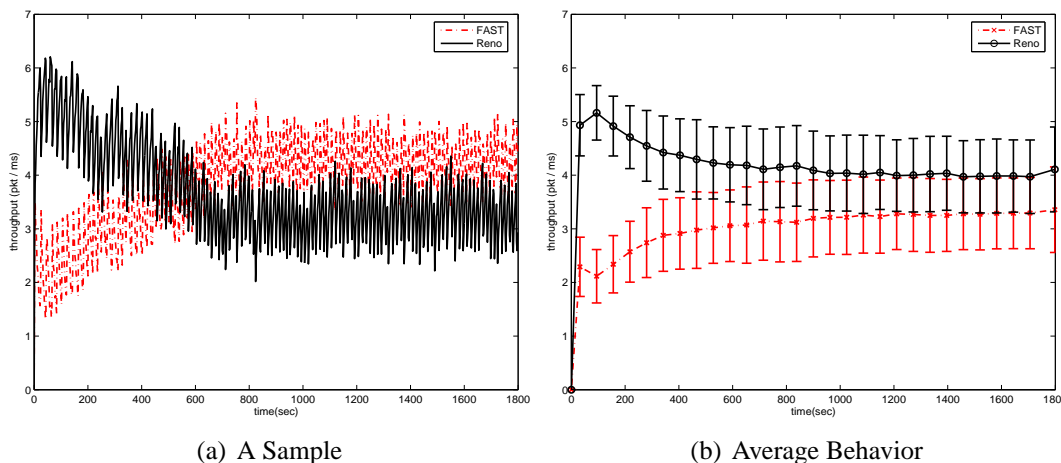


Figure 4.7: FAST vs. Reno, with buffer size of 400 pkts and Algorithm 1

The trajectory of α is presented in Figure 4.9. It is clear that although both starting for $\alpha = 50$, FAST finally uses a much larger α to deal with $B = 400$ case than $B = 80$ case as it experiences larger delay when $B = 400$.

4.2.2 Independence of bandwidth allocation on flow arrival pattern

Experiment 4.2b: independence of bandwidth allocation on flow arrival pattern

We repeat the simulations in Experiment 4.2a with Algorithm 1, w is set to be 1820s. Figure 4.10 and Figure 4.11 show the effect of α adaptation in the multiple-bottleneck case that we introduced in Example 2. As we proved in Theorem 4.1, there is always a unique equilibrium if we adapt α according to Algorithm 1. In this particular case, this single equilibrium is around the point where each Reno flow gets a throughput of 0.6 pkt per ms

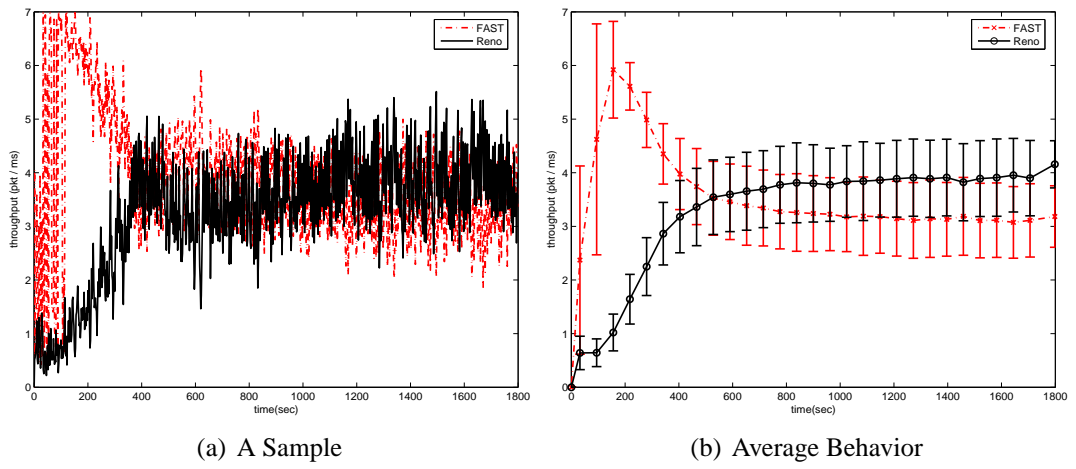


Figure 4.8: FAST vs. Reno, with buffer size of 80 pkts and Algorithm 1




Figure 4.9: α trajectory in experiment 4.1b

and each FAST flow gets 1.5 pkt per ms. At this single equilibrium, link 1 and link 3 are the bottleneck links. In Figure 4.10, FAST flows start on time zero and link 2 becomes the bottleneck. When Reno flows join on the 100th second, the ratio of queue to loss on link 2 is much higher than the target value. The FAST flows hence reduce their α values gradually and the bottleneck switches from link 2 to link 1 and 3 on 2000th second. After that, FAST flows and Reno flows converge to the unique equilibrium. The trajectory of α is presented in Figure 4.12. As we can see, depending on which flows start first, α follows a very different path although it finally reaches the same targeted value.

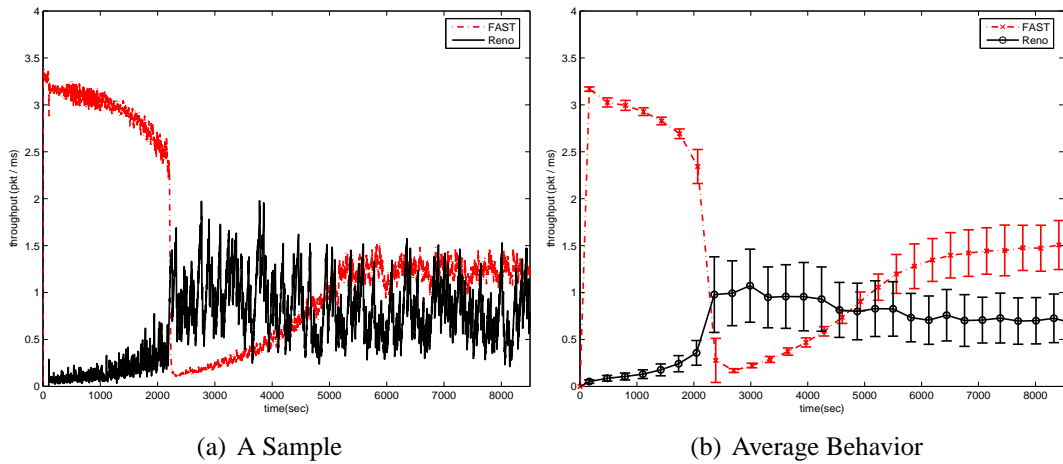


Figure 4.10: FAST starts first with Algorithm 1.

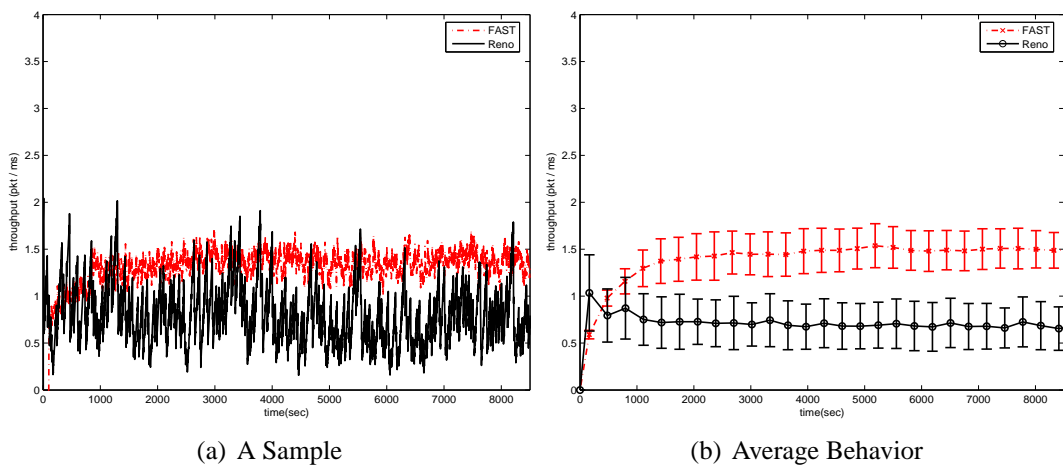
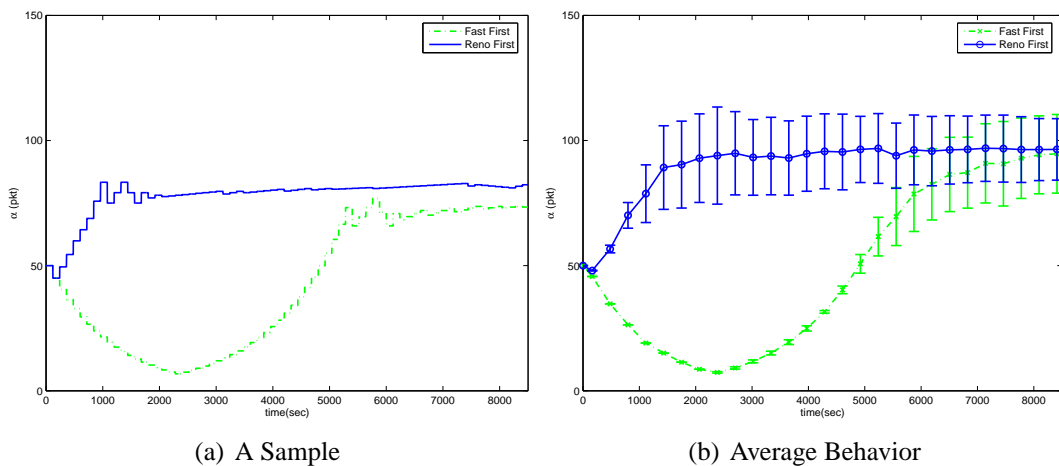


Figure 4.11: Reno starts first with Algorithm 1.

Figure 4.12: α trajectory in experiment 4.2b

4.3 Analysis

As pointed out in Corollary 3.2, all equilibria are Pareto efficient. However, no guarantee on fairness can be provided. We now turn from analysis to design, and develop a readily implementable control mechanism that “drives” any network with heterogeneous congestion control protocols to any desired operating point with a fair and efficient bandwidth allocation. This also explains the intuition behind and the theoretical foundation of Algorithm 1 in section 4.2. The central problems that motivate our study include: What is the equilibrium the system should be driven to? Can we make it unique? Will it solve any global optimization problem? How to do that in a distributed way? In this section, we propose an answer by introducing slow timescale updating. Our target equilibrium is still the maximizer of some weighted aggregate utility. The first step is to show the existence and uniqueness of such a solution.

Theorem 4.1. *For any given network (c, m, U, R) , for any positive vector w , there exists a unique positive vector μ such that if every source scales their own prices by μ_i^j , i.e.,*

$$x_i^j = (U_i^j)^{t-1} \left(\frac{1}{\mu_i^j} \sum m_l^j(p_l) \right)$$

then, at equilibrium (x^*, p^*) , x^* solves

$$\max_{x \geq 0} \quad \sum_{(i,j)} \frac{1}{w_i^j} U_i^j(x_i^j) \quad (4.2)$$

$$\text{subject to} \quad Rx \leq c \quad (4.3)$$

Moreover,

$$\mu_i^j = \frac{1}{w_i^j} \frac{\sum_{l \in L(j,i)} m_l^j(p_l^*)}{\sum_{l \in L(j,i)} p_l^*}$$

Proof. We claim that the optimality conditions of (4.2) and (4.3) are the same as equations that characterize the equilibrium of the above system ((2.2), (4.2), (2.4) and (2.5)). Capacity constraints, nonnegativity, and complementary slackness are obviously the same. We only need to check the relation between rates and prices at equilibrium. For our system, that are

$$\mu_i^j (U_i^j)'(x_i^j) = \sum_{l \in L(j,i)} m_l^j(p_l) \quad (4.4)$$

and

$$\mu_i^j = \frac{1}{w_i^j} \frac{\sum_{l \in L(j,i)} m_l^j(p_l^*)}{\sum_{l \in L(j,i)} p_l^*} \quad (4.5)$$

Combining them, we get

$$\frac{1}{w_i^j} (U_i^j)'(x_i^j) = \sum_{l \in L(j,i)} p_l \quad (4.6)$$

which is the relation between x and p specified by the optimality conditions of problem (4.2)-(4.3). On the other hand, given x and p that satisfy (4.6), one can always define μ_i^j by (4.5), and (4.4) will also be satisfied. \square

Remarks:

1. Parameter w enables us to measure fairness and to achieve any desired fair bandwidth allocation. The above result generalizes a theorem in [72], which asserts that by properly choosing α parameters in FAST flows, essentially any desired fairness

Algorithm 2 Two timescale control scheme

1. Every source chooses its rate by

$$x_i^j(t) = (U')^{-1} \left(\frac{q_i^j(t)}{\mu_i^j(t)} \right);$$

2. Every source updates its μ_i^j by

$$\mu_i^j(t+T) = \mu_i^j(t) + \kappa_i^j \left(\frac{\sum_{l \in L(j,i)} m_l^j(p_l(t+T))}{\sum_{l \in L(j,i)} p_l(t+T)} - \mu_i^j(t) \right)$$

where κ_i^j is stepsize for flow (j, i) and T is long enough so that the fast timescale dynamics among x and p can reach steady state.

between FAST and Reno is possible.

2. We only need all sources to have access to *one* common price. For example, when FAST and Reno coexist, since FAST also has access to the congestion price of packet loss, if it updates its parameter taking into account of loss, Theorem 4.1 holds.

Theorem 4.1 naturally suggests Algorithm 2 as a two-timescale scheme to control the operating point of networks with heterogenous congestion control protocols. The behavior of Algorithm 2 will be demonstrated in Section 4.4 through numerical examples.

In the extreme case, if μ_i^j is also updated at the same timescale as x_i^j , then sources all react to p and the system will be globally asymptotically stable. The essential idea in Algorithm 2 is that by reacting to the same price in slow timescale, uniqueness and fairness of equilibrium is guaranteed in the long run. Yet the algorithm allows sources to react to their own effective prices $m_i^j(p_l(t))$ at the fast timescale. This flexibility is important in practice when, for example, the link prices p_l are loss probability that are hard to reliably estimate at the fast timescale. The slow timescale algorithm only updates a linear scalar, which is readily implementable, e.g., this corresponds to update a parameter α in FAST/Vegas. Indeed, if we specialize Algorithm 2 to FAST/Neno networks using loss as the common price p , we get Algorithm 1 in section 4.2.

We now use the following coupled dynamical system to model TCP networks with the above scheme and study its convergence:

$$x_i^j(q_i^j(t)) = (U_i^j)^{\prime-1} \left(\frac{1}{\mu_i^j} q_i^j(t) \right) \quad (4.7)$$

$$q_i^j(t) = \sum_l R_{li}^j m_l^j(p_l(t)) \quad (4.8)$$

$$\dot{p}_l(t) = y_l(p(t)) - c_l \quad (4.9)$$

$$\epsilon \dot{\mu}_i^j(t) = \frac{q_i^j}{\sum_{l \in L(i,j)} p_l} - \mu_i^j \quad (4.10)$$

When ϵ is zero, differential equation (4.10) reduces to an algebraic equation and μ_i^j is updated instantaneously. Sources essentially react to p and the system is globally asymptotically stable. The following theorem shows that globally asymptotic stability still holds when we introduce dynamics to update $\mu_i^j(t)$ as in (4.10).

Theorem 4.2. *There exists an $\epsilon^* > 0$, such that for all $0 < \epsilon < \epsilon^*$, the system described by (4.7)–(4.10) is globally asymptotically stable.*

Proof. See Appendix 6.7. □

4.4 Numerical Examples

Throughout this subsection, we provide some Matlab numerical results to further validate the effectiveness of the control scheme proposed in section 4.2. For simplicity we choose w to be a vector with all components being 1, i.e., we attempt to maximize the aggregate utility.

Example 4.1: L=2

We consider a two-link network with six flows using two prices. The routing matrices are:

$$R^1 = R^2 = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

All sources have log utility functions. Both links have capacities 10 units. We start with $\mu_i^j = 1$. In each step, we directly update $\mu_i^j = \frac{\sum m(p)}{\sum p}$ ($\kappa_i^j = 1$). We can easily find the equi-

librium by solving the utility maximization problem and obtain the following equilibrium,

$$x_1^1 = x_2^1 = x_1^2 = x_2^2 = 10/3, x_3^1 = x_3^2 = 5/3, p_1 = p_2 = 0.3.$$

Case 1: The price mapping functions are linear : $m_l^j(p_l) = k_l^j p_l$ where

$$K^1 = I, \quad K^2 = \text{diag}(2, 4)$$

We show the steady state prices after each iteration in Table 4.2. μ_2^3 for each step is also shown. Note that other μ_i^j 's are always 1 in this case. The system converges to the predicted equilibrium after three iterations.

Table 4.2: Steady state after each iteration: case 1

| Iterations | μ_3^2 | p_1 | p_2 |
|------------|-----------|--------|--------|
| 1 | 1.0000 | 0.2233 | 0.1861 |
| 2 | 2.9091 | 0.2985 | 0.2985 |
| 3 | 3.0000 | 0.3000 | 0.3000 |

We then move to the case of nonlinear price mapping function. In the following two cases, we again show the steady states after each iteration and its convergence to the target equilibrium. See Table 4.3 and 4.4.

Case 2: $m_1^2(p_1) = (p_1)^2, m_2^2(p_2) = 2(p_2)^2$.

Case 3: $m_1^2(p_1) = \sqrt{p_1}, m_2^2(p_2) = 3\sqrt{p_2}$.

Table 4.3: Steady state after each iteration: case 2

| Iterations | μ_1^2 | μ_2^2 | μ_3^2 | p_1 | p_2 |
|------------|-----------|-----------|-----------|--------|--------|
| 1 | 1.0000 | 1.0000 | 1.0000 | 0.4685 | 0.3582 |
| 2 | 0.4685 | 0.7163 | 0.5759 | 0.3489 | 0.3171 |
| 3 | 0.3489 | 0.6342 | 0.4848 | 0.3151 | 0.3050 |
| 4 | 0.3151 | 0.6100 | 0.4601 | 0.3047 | 0.3014 |
| 5 | 0.3047 | 0.6028 | 0.4530 | 0.3015 | 0.3004 |
| 6 | 0.3015 | 0.6008 | 0.4509 | 0.3005 | 0.3001 |

Example 4.2: L=3 with multiple equilibria

Table 4.4: Steady state after each iteration: case 3

| Iterations | μ_1^2 | μ_2^2 | μ_3^2 | p_1 | p_2 |
|------------|-----------|-----------|-----------|--------|--------|
| 1 | 1.0000 | 1.0000 | 1.0000 | 0.2155 | 0.1672 |
| 2 | 2.1543 | 7.3364 | 4.4187 | 0.3331 | 0.3691 |
| 3 | 1.7326 | 4.9380 | 3.4174 | 0.2909 | 0.2817 |
| 4 | 1.8542 | 5.6520 | 3.7228 | 0.3208 | 0.3061 |
| 5 | 1.8174 | 5.4228 | 3.6299 | 0.2992 | 0.2981 |
| 6 | 1.8282 | 5.4943 | 3.6580 | 0.3002 | 0.3006 |

In this experiment, we use the following example that was used in example 2.3 (section 2.3.2) to demonstrate existence of multiple isolated equilibria. The network is shown in Figure 2.1 with three unit-capacity links, $c_l = 1$. There are three different protocols with the corresponding routing matrices

$$R^1 = I, \quad R^2 = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}^T, \quad R^3 = (1, 1, 1)^T$$

The price mapping functions are assumed to be linear with coefficients

$$K^1 = I, \quad K^2 = \text{diag}(5, 1, 5), \quad K^3 = \text{diag}(1, 3, 1)$$

Utility functions of sources (j, i) are

$$U_i^j(x_i^j, \alpha_i^j) = \begin{cases} \beta_i^j (x_i^j)^{1-\alpha_i^j} / (1 - \alpha_i^j) & \text{if } \alpha_i^j \neq 1 \\ \beta_i^j \log x_i^j & \text{if } \alpha_i^j = 1 \end{cases}$$

with appropriately chosen positive constants α_i^j and β_i^j shown in table 2.1 (section 2.3.2). These utility functions can be viewed as a weighted version of the α -fairness utility functions proposed in [55]. μ_i^j 's are updated every 20 time units. We show that although the system reaches different equilibria after the first iteration, it nevertheless finally reaches the unique target. In terms of convergence time, not surprisingly, both being too cautious ($\kappa_i^j = 0.1$) and too aggressive ($\kappa_i^j = 0.9$) are not optimal, which can be clearly seen by comparing with the $\kappa_i^j = 0.5$ case.

Case 1: We start with initial point $p_1(0) = p_2(0) = p_3(0) = 0.3$. After the first iteration,

the network goes to equilibrium ($p_1^* = p_3^* = 0.165$, $p_2^* = 0.170$). $p_1(t)$ with different updating stepsize κ_i^j is shown in Figures 4.13.




Figure 4.13: Case 1: $p_1(t)$ with different κ_i^j

Case 2: We choose another initial point $p_1(0) = p_3(0) = 0.1$, $p_2(0) = 0.3$ As shown in Figure.4.14. After the first iteration, the system reaches another equilibrium, $p_1^* = p_3^* = 0.135$ and $p_2^* = 0.230$. However finally, the system still reaches the same steady state as in Figure 4.14 (Both converge to about 0.222).




Figure 4.14: Case 2: $p_1(t)$ with different κ_i^j

Example 4.3: L=5 with asynchronous updating

In this experiment, we have a larger network with five links and 15 flows. Also, the scheme is tested in an asynchronous environment. We assume every five time units, flows can update their μ_i^j and they do so with some probability. Hence every five time units, only a portion of flows update their μ_i^j .

We randomly set link capacities uniformly between 1 and 10, and take price mapping functions to be $m^1(p) = p$ and $m^2(p) = p^\alpha$, where α is randomly chosen between 0.5 and 5 with uniform distribution. Flows 1 to 5 use links 1 to 5 respectively while a random routing matrix with entries 0 or 1 with equal probability is used to define routes for other flows. Finally each flow randomly chooses to use price 1 or 2 with equal probability.

All of the 1000 trials converge to the right target. Some typical convergence patterns are shown in Figure 4.15. It shows clearly that although asynchronism causes longer convergence time, the system still converges to the same equilibrium.

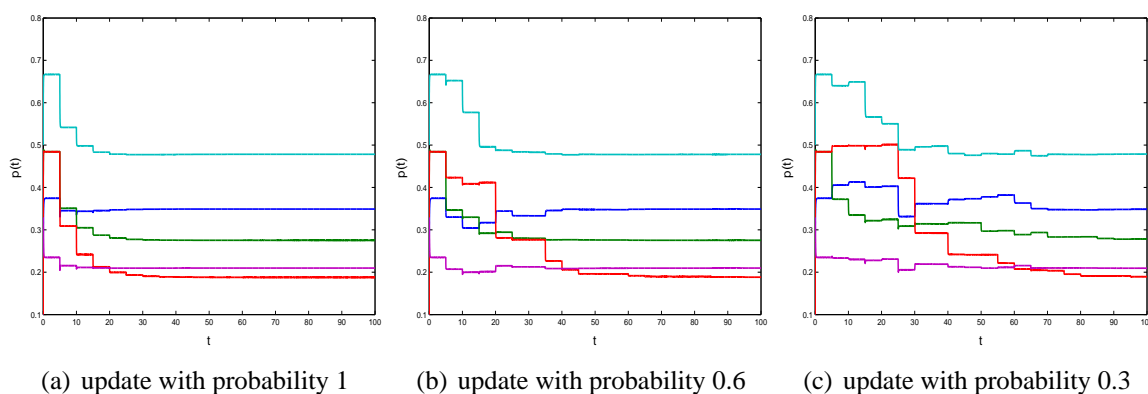


Figure 4.15: $p(t)$ with different probability of updating

4.5 WAN in Lab Experiments

The objective of experiments in this subsection is to show the effectiveness of our slow timescale updating algorithm in a more realistic setting. We achieve that by carrying out experiments with TCP Reno and FAST in WAN in Lab [2] and by considering more practical scenarios not considered before (e.g., small buffer size, only FAST flows).

WAN in Lab is a wide area network consisting of an array of reconfigurable routers, servers and clients. The backbone of the network is connected by two 1600 km OC-48 links that can provide a large real propagation delay. We test our algorithm with a single bottleneck link shown in Figure 4.16.

Experiment 4.3: small buffer size

WAN-in-LAB Testbed

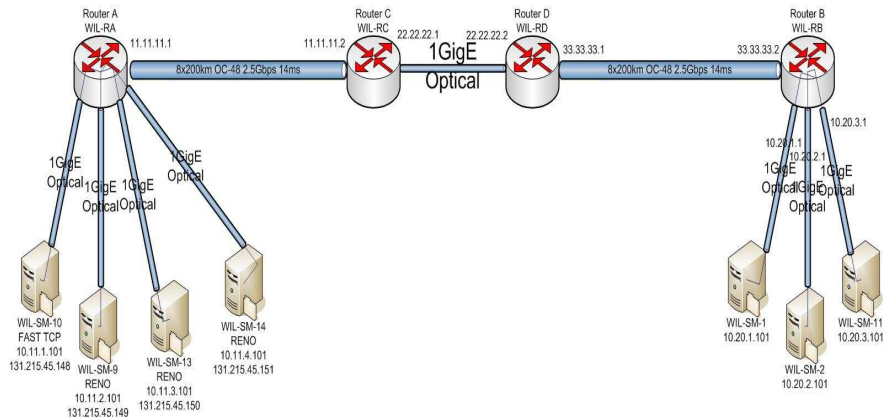


Figure 4.16: WAN in Lab experiment setup

In this experiment, we show a byproduct of our slow timescale algorithm, namely automatically matching FAST's parameter α to the buffer size in the network. As we know, every FAST flow tries to maintain α packets in the queues along its path. Clearly if the buffer capacity is smaller than the α used, a constant high packet loss rate will occur and both Reno and FAST will have very poor throughput. Here we show that our algorithm can adjust α automatically to a proper value when it sees a high loss rate.

One FAST and one Reno compete for bandwidth of the bottleneck link with 1Gbps (80pkts/ms) capacity. The buffer capacity is 480pkts. The initial α is set to be $\alpha_0=800$. The results are summarized in Figure 4.17. As the left part of the figure shows, both Reno and FAST get very low throughput due to the high packet loss rate (FAST: 135Mbps; Reno: 22Mbps). However, using the slow timescale update, FAST decreases its α as it sees high loss and finally both flows get high throughput (FAST: 593Mbps; Reno: 246Mbps). The utilization is increased dramatically from 15.7 percent to 83.9 percent.

Experiment 4.4: only FAST flows

Although the slow timescale update shows desirable properties in various tests we have discussed so far, there is a problem we have not touched, namely the case when there are only FAST flows in a network. As FAST is designed to achieve a steady state with no loss, flows will keep increasing their α until the buffer is filled and loss is generated. This is not

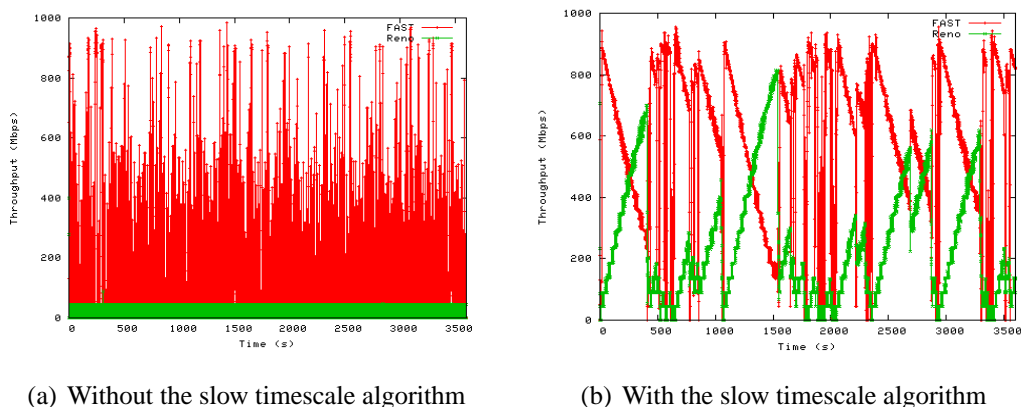


Figure 4.17: Bandwidth partition between Reno and FAST

desirable and we propose to turn off the slow timescale updating algorithm when a flow has not seen any loss for a certain amount of time (one second by default). We conduct an experiment using three FAST flows all with $\alpha_0=200$ and test this idea. The throughput trajectory is shown in Figure 4.18. We can see that after a period of adjusting, all flows are stabilized. The steady state throughputs are 128Mbps, 234Mbps and 566Mbps, which result in a high utilization of 92.8 percent even though the initial sum of α exceeds the buffer capacity. However, this introduces potential fairness problem as we cannot control the exact α values when they stop updating when no loss is generated. For example, instead of achieving perfect fairness with a Jain index [34] of 1, we have 0.733 in this experiment. We tend to think that this short term unfairness is not so important as in practice, flows come and go, which will give many chances for existing flows to reshuffle and the random short term unfairness can be averaged out to yield long term fairness.




Figure 4.18: Bandwidth sharing among FAST flows

Chapter 5

Conclusion

Writing is a deed of eternity. Its ups and downs are known only to the author.

— Du Fu (712 – 770), A Great Poet in Tang Dynasty, China

The central motivation of this thesis is to investigate end-to-end flow control in an environment with multiple pricing signals, or more generally, the resource allocation problem with heterogeneous prices. It is demonstrated in this thesis that when sources sharing the same network react to different pricing signals, the current duality model no longer explains the equilibrium of bandwidth allocation. We have introduced a mathematical formulation of network equilibrium for multi-protocol networks and studied several fundamental properties such as existence, uniqueness, optimality, and stability. We prove that equilibria exist, and are almost always locally unique. The number of equilibria is almost always finite and must be odd. The equilibrium is globally unique if the price mapping functions are similar, or the $J(p)$ is “negative definite” along certain directions. By identifying an optimization problem associated with *every* equilibrium, we show that every equilibrium is Pareto efficient. It also yields an upper bound on efficiency loss due to pricing heterogeneity. On fairness, we show that intra-protocol fairness is still decided by a utility maximization problem while inter-protocol fairness is the part over which we don’t have control. However it is shown that we can achieve any desirable inter-protocol fairness by properly choosing protocol parameters. Regarding dynamics, various stability results are provided. In particular we prove that if the degree of pricing heterogeneity is properly bounded, then the network equilibrium is not only unique but also locally stable. Finally, we propose a scheme to steer

an arbitrary network to a unique equilibrium that maximizes the total utility, by updating in a slow timescale a linear parameter in sources' algorithms. The scheme uses only local information. In addition to analysis and numerical examples, we have presented NS-2 simulations and WAN in Lab experiments using TCP Reno and FAST to demonstrate the correctness and convergence of the scheme.

There are a number of features of this study. First, our emphasis is on general networks with multiple sources and links that use a large class of algorithms to adapt their rates and congestion prices. Often, interesting and counter-intuitive behaviors arise only in a network setting where sources interact through shared links in intricate and surprising ways [73, 74]. Such behaviors are absent in single-link models and are usually hard to discover or explain without a fundamental understanding of the underlying structure. Second, starting from a concrete engineering system, we set up a mathematical framework to further explore structures, clarify ideas, and suggest improvements. In the process of doing so, we have borrowed some tools and techniques from other communities that are not widely used in the field of communication networking, e.g., general equilibrium analysis from economics, and Sard's theorem and Poincare-Hopf index theorem from differential geometry. More interestingly, we have also developed some new mathematical techniques; see Appendix 6.2 and 6.6 for example. These results can potentially be useful in solving important problems in economics and mathematics, as we will discuss in more detail later. Finally, though it is mainly theoretical, this thesis also includes experimental verification of its key predictions. These supporting data come from a range of methods that span from numerical calculation and packet-level simulations to Dummynet and WAN in Lab experiments.

It is always exciting to look ahead. We now conclude by discussing possible future directions. There are two natural extensions we need to address.

The first question is that of the global stability of heterogeneous congestion control protocols. We know there can be multiple equilibria, and they cannot all be locally stable unless there is only one. We have conditions under which the equilibrium is unique and locally stable. However, we still don't know the global dynamics of the system in general. One plausible conjecture is that every trajectory ends up with one of the equilibria. The intuition is that heterogeneity of prices makes the underlying dynamics no longer an exact

gradient field, but the monotonicity of price mapping functions still guarantees that the system reacts in a qualitatively correct way. In general, mathematically showing this is quite hard as only limited tools are available to deal with multiple equilibria, e.g., recent developments in monotone dynamic systems [69] and the dual of the Lyapunov method [60].

The other problem concerns the convergence property of the slow timescale update scheme. We have shown that, if the update is fast enough, the system will converge to the unique optimal point. However, in reality, we want to update much more slowly compared to the underlying TCP dynamics. Addressing global stability for this hybrid system then becomes very challenging, especially given that the underlying TCP system can have multiple equilibria. The current conjecture is that if the stepsize sequence satisfies some certain condition [11], the system converges globally.

This thesis deals with a problem at the intersection of engineering and economics and also involves mathematics such as asymmetric matrix and vector field analysis. Some results and techniques here can potentially have much wider applications beyond the study of congestion control systems. We now list a few long term directions, one for each discipline.

Economics: “Bounded heterogeneity implies regularity” Economists began to seriously study dynamics of market behaviors long time ago when Hicks advocated its importance in [30], in which he proposed his stability concepts (imperfect stability and perfect stability). It was later found that Hicksian stability is only remotely related to real dynamic stability, which is pointed out by Samuelson [63].

In general equilibrium theory in economics, with the classical Arrow-Debreu model (see section 2.4 for the version of pure exchange economics and [6, 51] for more details), existence, uniqueness, and stability¹ of equilibrium have been carefully and rigorously investigated in the last half century. Unlike existence, which was a triumph of mathematical economics [19], results on global uniqueness and stability are limited in the sense that the conditions to guarantee those properties are quite restrictive. Moreover, failure to provide applicable results on global uniqueness and stability weakens the significance of existence

¹The most commonly studied stability model is the Tatonnement process, whose differential equation version was proposed in [64]. One is referred to [5, 7, 8, 65, 68, 13, 24, 20, 28] for the development.

and welfare theorems. These all cause confusion and even skepticism about the future of general equilibrium theory [50, 4].

We observe that the examples that have multiple equilibria or limit cycles typically involve some assumptions that may not accurately reflect reality. For example, in his classical work [65], Scarf demonstrated that it is possible for a pure exchange economy with three consumers to have a globally unique equilibrium that is not even locally stable. The utility functions of users are completely complementary in the sense that consumer 1 values commodity A while commodities B and C do not affect its utility, while consumers 2 and 3 value commodities B and C, respectively.

A natural question is then whether we can provide a bound on heterogeneity of consumers' utility functions or their initial endowments to guarantee a unique and globally attractive equilibrium. In short, does "bounded heterogeneity imply regularity"? By regularity, we mean desirable properties such as uniqueness and stability. If that turns out to be true, then we can potentially overcome the difficulty of general equilibrium theory by justifying the needed bound on the degree of heterogeneity using some statistical argument for large systems.

Engineering: "How to use multiple prices optimally" We have analyzed networks with heterogeneous congestion control protocols where every protocol uses exactly one kind of price. As sources may have access to multiple prices, it is interesting to consider the optimal way for sources to regulate their rates based on all information they can observe. Some steps are taken in this direction by combining delay-based and loss-based congestion control protocols [15, 71]. We are interested in a general framework under which we can explore various basic questions, e.g., Is there any performance limit due to the finite feedback information? If so, how do we express that and what is its implication in practice? Hopefully the answers to these questions can lead to more systematic designs.

Mathematics: "Asymmetric P -matrix analysis" The general matrix stability problem has been studied for 200 years. See [26] for detailed classic results and [29] for an excellent survey on recent progress. In the last 50 years, people have begun to look at nonsymmetric cases. We take the P -matrix as an example here.

A P -matrix is a real square matrix all of whose principal minors are positive. Positive

definite matrices are symmetric P -matrices. One of the important questions being asked in the linear algebra community is, what is the required “symmetry” to guarantee the stability of a P -matrix, i.e., all its eigenvalues have positive real parts. This is a very difficult problem and so far the only general (but conservative) result is due to Carlson [14], which asserts that every positive sign-symmetric matrix is positive stable. A matrix A is said to be positive sign-symmetric if it is a P -matrix and $A(\alpha, \beta)A(\beta, \alpha) \geq 0$ for all $\alpha, \beta \in Q_n$, $|\alpha| = |\beta|$. Here, $Q_n = \{(i_1, i_2, \dots, i_k) \mid 1 \leq i_1 < i_2 < \dots < i_k \leq n\}$. Given a matrix A and $\alpha, \beta \in Q_n$, $A(\alpha, \beta)$ means the minor of A whose rows are indexed by α and whose columns are indexed by β .

People have tried to identify a set of properties from positive definite matrices that is critical for the stability of P -matrices. Various conjectures have been proposed, but all fail for the general case. One can find these results in [31] and references therein.

Motivated by the analysis of $\mathbf{J}(p)$ in this thesis, which is a P -matrix once the price similarity condition holds, we have also carried out study on general P -matrices. In particular, we were able to find another condition on the degree of symmetry that is sufficient for the stability of a P -matrix [3]. A current conjecture is that a P -matrix is stable if its diagonal terms are all positive and larger than the absolute value of any term that is in the same row (column). Clearly, this is a stronger statement than Theorem 3.8. Theorem 3.6’s proof covers cases when the dimension of the matrix does not exceed three. It remains to see to what extent techniques in section 6.6 can be used to deal with this problem.

Chapter 6

Appendix

6.1 Simulation of Multiple Equilibria

The Dummynet experiments provide qualitative evidence of multiple equilibria with practical protocols. We could not have verified the experimental results with quantitative predictions because a Droptail router does not admit an accurate mathematical model for the price mapping function m_l . In this section, we present simulation results using NS-2 on multiple equilibria and fairness. The network simulator NS-2 allows us to use RED routers for which the price mapping function m_l is known. We can thus compare simulation measurements with our theoretical predictions. For all the simulations in this section, TCP Vegas is used, which has the same equilibrium structure as FAST.

The network simulator ns-2 version 2.1b9a is used here. We use the RED algorithm and packet marking instead of dropping. The marking probability $p(b)$ can be expressed as in (4.1).

The network topology is as shown in Figure 1.2. The link capacities of link 1 and link 3 are set to be 100 Mbps (8.33pkts/ms) and the one way propagation delay to be 50 ms. For link 2, the capacity is 150 Mbps (12.5pkts/ms) and one way propagation delay is 5 ms. There are 10 Vegas flows on each of paths 1 and 2, and 20 Reno flows on path 3. As in NS simulations, αd is the number of packets the flow maintains along its path, which has been called α before by convention. Hence every flow tries to put 5.5 packets along its path as we set $\alpha = 50$.

Experiment 6.1: varying K_2 .

We set $(\underline{b}_1, \bar{b}_1, K_1)$ to be $(0, 1000, 10000)$ at link 1 and link 3. Set $(\underline{b}_2, \bar{b}_2)$ to be $(100, 1500)$ at link 2, and vary the slope K_2 at link 2 from 10 to 600. Figure 6.1 shows the aggregate throughput of all Reno flows and the link utilization at link 1 for different values of K_2 . Theoretical predictions are calculated by solving equilibrium equations and the price mapping function (4.1) for RED. As can be seen, the prediction matches the measured curve very well.

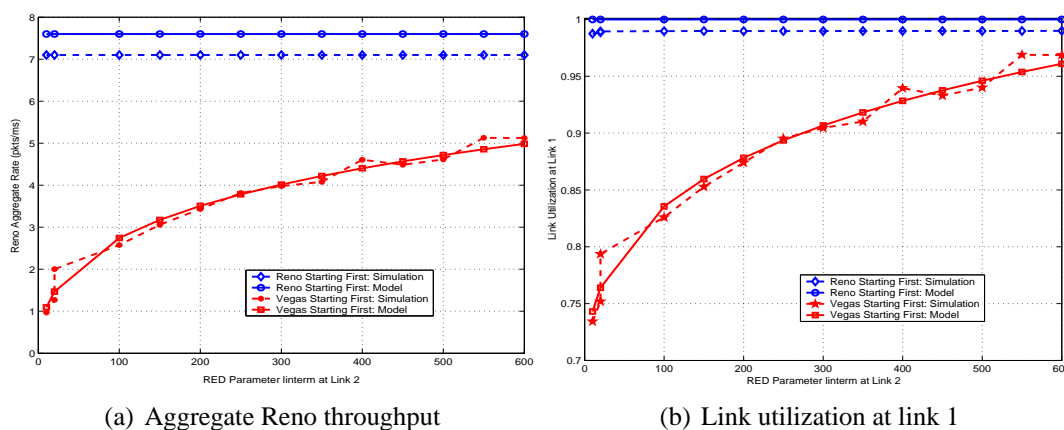


Figure 6.1: Experiment 6.1: Aggregate Reno throughput and link utilization at link 1.

From Figure 6.1, the aggregate throughput and utilization at link 1 are independent of K_2 if Reno flows start first. This is because link 2 is not saturated in this scenario, as explained earlier, and hence varying its parameter does not affect the equilibrium. When Vegas flows start first, on the other hand, link 2 is the bottleneck link, and hence as K_2 increases, Reno achieves more and more bandwidth since the mapping function penalizes Reno less and less.

As K_2 increases, one may expect that the Reno throughput curve in Figure 6.1 that correspond to Vegas starting first will converge to the same value for the case when Reno starts first. It is not possible to exhibit this beyond $K_2 = 600$ at link 2. As shown in Figure 6.1, the utilization at link 1 is more than 95% when $K_2 = 600$. Even though link 1 is not saturated yet, it is so close to being saturated that random fluctuations in the queue can readily shift the system from the current equilibrium where only link 2 is saturated to the other equilibrium where links 1 and 3 are saturated (while link 2 is not). See a clear

demonstration of this phenomenon in Experiment 6.3.

Experiment 6.2: varying K_1 .

In this experiment, we fix $K_2 = 100$ at link 2 and vary K_1 at link 1 and link 3 simultaneously from 5,000 to 11,000. The results are summarized in Figure 6.2. When Vegas

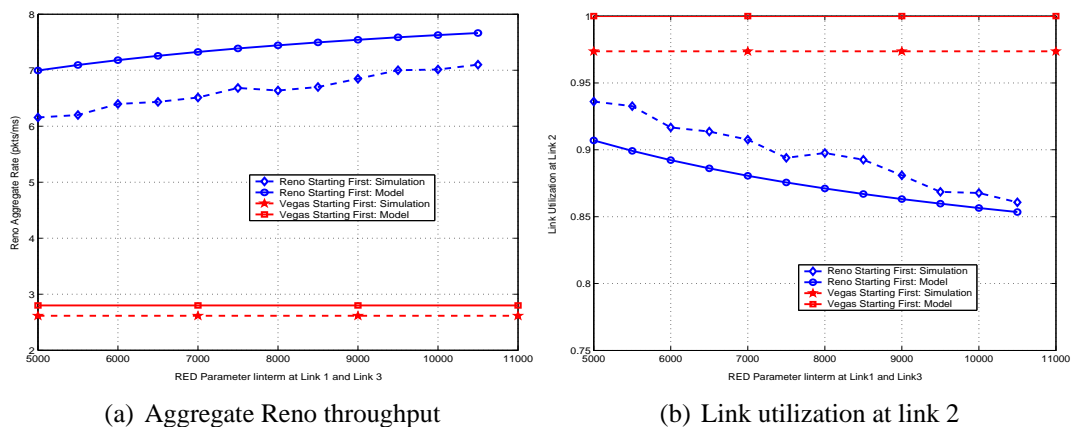


Figure 6.2: Experiment 6.2: Aggregate Reno throughput and link utilization at link 2.

flows start first, the bottleneck link is link 2 and therefore both the aggregate Reno throughput and the utilization at link 2 are independent of K_1 . When Reno flows start first, on the other hand, links 1 and 3 become saturated and varying K_1 affect both the aggregate Reno throughput and link 2's utilization. The theoretical predictions track the measured data, but are generally larger than the data. The main reason is that Vegas flows overestimated base RTT when Reno flows start first and maintain a nonzero queue. Then Vegas flows become more aggressive and suppress Reno flows more than they should; see [47] for more discussion on the effect of error in base RTT estimation.

As K_1 decreases at links 1 and 3, Reno flows see more losses and the system may shift to the other equilibrium where only link 2 is saturated. For instance, from Figure 6.2, the utilization at link 2 is close to 95 percent when $K_1 = 5000$.

Experiment 6.3: shifting equilibria.

This experiment shows that the system can shift back and forth between the two equilibria when the utilization of the unsaturated link(s) is sufficiently close to 100 percent so that the system can readily jump between two disjoint active constraint sets due to random

fluctuation. The slopes $K_1 = 3500$ at link 1 and link 3 and $K_2 = 500$ at link 2. The simulation duration is 1000 sec. The queues at link 1 and link 2 are shown in Figure 6.3. This result unambiguously exhibits that there are two equilibria and they are both achieved.




Figure 6.3: Experiment 6.3: Queue sizes at link 1 and link 2. The system shifts between the two equilibria with disjoint active constraint sets.

6.2 Proof of Theorem 2.8

By Corollary 2.7, we only need to prove that $I(p) = (-1)^L$ for any equilibrium $p \in E$. Since $\det(\mathbf{J}(p)) = (-1)^L \det(-\mathbf{J}(p))$, the condition reduces to $\det(-\mathbf{J}(p)) > 0$. Now

$$\begin{aligned} -\mathbf{J}(p) &= -\sum_j R^j D^j(p) (R^j)^T \frac{\partial m^j}{\partial p}(p) \\ &= \sum_j B^j M^j \end{aligned}$$

where $M^j = M^j(p) = \frac{\partial m^j}{\partial p}(p)$ is a diagonal matrix, and $B^j = B^j(p)$ is defined by its elements

$$B_{kl}^j = \sum_i R_{ki} R_{li} \left(-\frac{\partial^2 U_i^j}{\partial (x_i^j)^2} \right)^{-1} \quad (6.1)$$

Hence

$$\begin{aligned} \det(-\mathbf{J}(p)) &= \det \left[\sum_j B^j M^j \right] \\ &= \sum_{\mathbf{k}} \text{sgn} \mathbf{k} \prod_{l=1}^L \sum_{j=1}^J [B^j M^j]_{k_l l} \end{aligned} \quad (6.2)$$

Here, the summation over $\mathbf{k} = (k_1, \dots, k_L) \in \{1, \dots, L\}^L$ is over all $L!$ permutations of the L items $\{1, \dots, L\}$. The function $\text{sgn} \mathbf{k}$ is 1 if the minimum number of pairwise interchanges necessary to achieve the permutation \mathbf{k} starting from $(1, 2, \dots, L)$ is even and -1 if it is odd.

Let π denote an L -bit binary sequence that represents the path consisting of exactly those links k for which the k th entries of π are 1, i.e., $\pi_k = 1$. Let $\Pi(k, l) := \{\pi | \pi_k = \pi_l = 1\}$ be the set of paths that contain both links k and l . Let $I_\pi^j = \{i | R_{l_i}^j = 1 \text{ if and only if } \pi_l = 1\}$ be the set of type j sources on path π , possibly empty. Let

$$r_\pi^j = r_\pi^j(p) = \sum_{i \in I_\pi^j} \left(-\frac{\partial^2 U_i^j}{\partial (x_i^j)^2} \right)^{-1} \quad (6.3)$$

where r_π^j is zero if I_π^j is empty. Since all utility functions are assumed concave, $r_\pi^j \geq 0$. Then we have from (6.1) and (6.3)

$$B_{kl}^j = \sum_{\pi \in \Pi(k, l)} r_\pi^j \quad (6.4)$$

This together with (6.2) implies

$$\det(-\mathbf{J}(p)) = \sum_{\mathbf{k}} \text{sgn} \mathbf{k} \prod_{l=1}^L \sum_{j=1}^J \left(m_l^j \sum_{\pi \in \Pi(k_l, l)} r_\pi^j \right) \quad (6.5)$$

Consider any sequence $a_{ij}, j \in J_i, i = 1, \dots, I$, where J_i is a finite index set that

depends on i . We have

$$\prod_{i=1}^I \sum_{j \in J_i} a_{ij} = \sum_{\mathbf{j}} \prod_{i=1}^I a_{ij_i} \quad (6.6)$$

where \mathbf{j} denotes the vector index $\mathbf{j} = (j_1, \dots, j_I)$ and the summation is over all values in $J_1 \times \dots \times J_I$.

Using (6.6) to change the order of product over l and summation over j in (6.5), we have

$$\det(-\mathbf{J}(p)) = \sum_{\mathbf{k}} \text{sgn} \mathbf{k} \sum_{\mathbf{j}} \prod_{l=1}^L \left(\dot{m}_l^{j_l} \sum_{\pi \in \Pi(k_l, l)} r_{\pi}^{j_l} \right)$$

where the vector index $\mathbf{j} = (j_1, \dots, j_L)$ ranges over $\{1, \dots, J\}^L$. Applying (6.6) again to change the order of product over l and summation over the index π , we have

$$\det(-\mathbf{J}(p)) = \sum_{\mathbf{k}} \text{sgn} \mathbf{k} \sum_{\mathbf{j}} \mu(\mathbf{j}) \sum_{\pi \in \Pi(\mathbf{k}, \mathbf{l})} \rho(\mathbf{j}, \pi) \quad (6.7)$$

where

$$\mu(\mathbf{j}) := \prod_{l=1}^L \dot{m}_l^{j_l} \quad (6.8)$$

$$\rho(\mathbf{j}, \pi) := \prod_{l=1}^L r_{\pi^l}^{j_l} \quad (6.9)$$

The last summation in (6.7) is over the vector index $\pi = (\pi^1, \dots, \pi^L)$ that takes value in the set $\{\text{all } L\text{-bit binary sequences}\}^L$. As mentioned above, $\mathbf{l} = (1, \dots, L)$ denotes the identity permutation, and “ $\pi \in \Pi(\mathbf{k}, \mathbf{l})$ ” is a shorthand for “ $\pi^l \in \Pi(k_l, l), l = 1, \dots, L$ ”. Denote by $\mathbf{1}(a)$ the indicator function that is 1 if the assertion a is true and 0 otherwise. Then (6.7) becomes

$$\det(-\mathbf{J}(p)) = \sum_{\mathbf{j}} \sum_{\pi} C(\mathbf{j}, \pi) \rho(\mathbf{j}, \pi) \quad (6.10)$$

where

$$C(\mathbf{j}, \boldsymbol{\pi}) := \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \mathbf{l})) \operatorname{sgn} \mathbf{k} \mu(\mathbf{j}) \quad (6.11)$$

Hence $\det(-\mathbf{J}(p))$ is a summation, over the index $(\mathbf{j}, \boldsymbol{\pi})$, of terms $\rho(\mathbf{j}, \boldsymbol{\pi})$ with coefficients $C(\mathbf{j}, \boldsymbol{\pi})$. We now show that only those terms for which the constituent r_{π}^j in the product $\rho(\mathbf{j}, \boldsymbol{\pi})$ are all distinct have nonzero coefficients.

Lemma 6.1. *Consider a term in the summation in (6.21) indexed by $(\mathbf{j}, \boldsymbol{\pi})$. If there are integers $a, b \in \{1, \dots, L\}$ such that $j_a = j_b$ and $\pi^a = \pi^b$, then $C(\mathbf{j}, \boldsymbol{\pi}) = 0$.*

Proof. Fix any $(\mathbf{j}, \boldsymbol{\pi})$. Suppose without loss of generality that $j_1 = j_2$ and $\pi^1 = \pi^2$ and $\rho(\mathbf{j}, \boldsymbol{\pi}) \neq 0$. We now show that its coefficient $C(\mathbf{j}, \boldsymbol{\pi}) = 0$.

Consider any permutation \mathbf{k} in (6.11) that gives a nonzero coefficient in $C(\mathbf{j}, \boldsymbol{\pi})$:

$$\mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \mathbf{l})) \operatorname{sgn} \mathbf{k} \mu(\mathbf{j}) = \operatorname{sgn} \mathbf{k} \mu(\mathbf{j}) \quad (6.12)$$

This means that

$$\pi^1 \in \Pi(k_1, 1) \quad \text{and} \quad \pi^2 \in \Pi(k_2, 2)$$

Hence, since $\pi^1 = \pi^2$, the path π^1 goes through all links $1, 2, k_1, k_2$. In particular

$$\pi^1 \in \Pi(k_2, 1) \quad \text{and} \quad \pi^2 \in \Pi(k_1, 2)$$

Therefore there is a permutation $\hat{\mathbf{k}}$ in (6.11) with $\hat{k}_1 = k_2$, $\hat{k}_2 = k_1$, and $\hat{k}_l = k_l$ for $l \geq 3$ for which $\mathbf{1}(\boldsymbol{\pi} \in \Pi(\hat{\mathbf{k}}, \mathbf{l})) = 1$ but $\operatorname{sgn} \hat{\mathbf{k}} = -\operatorname{sgn} \mathbf{k}$. This yields a term $-\operatorname{sgn} \mathbf{k} \mu(\mathbf{j})$ in $C(\mathbf{j}, \boldsymbol{\pi})$ which exactly cancels the term in (6.12). Since the argument applies to any \mathbf{k} in (6.11), $C(\mathbf{j}, \boldsymbol{\pi}) = 0$. \square

In view of Lemma 6.1, we will restrict the summation over the index $(\mathbf{j}, \boldsymbol{\pi})$ in (6.21) to the largest subset of $\{1, \dots, J\}^L$ where the constituent r_{π}^j in $\rho(\mathbf{j}, \boldsymbol{\pi})$ are all distinct. Let Θ

denote this subset. We abuse notation and define permutation $\sigma \in \{1, \dots, L\}^L$ on Θ by

$$\sigma(\mathbf{j}, \boldsymbol{\pi}) = (\sigma(\mathbf{j}), \sigma(\boldsymbol{\pi}))$$

Then let Θ_0 be the largest subset of Θ that is *permutationally distinct*, i.e., no vector in Θ_0 is a permutation of another vector in Θ_0 . The set of permutations $\sigma \in \{1, \dots, L\}^L$ is in one-one correspondence with the set of $(\mathbf{j}', \boldsymbol{\pi}')$ that are permutations of a given $(\mathbf{j}, \boldsymbol{\pi})$ in Θ_0 .¹ This allows us to carry out the summation over $(\mathbf{j}, \boldsymbol{\pi})$ in (6.21) first over $(\mathbf{j}, \boldsymbol{\pi})$ that are permutationally distinct and then over all their permutations. Notice that, given any $(\mathbf{j}, \boldsymbol{\pi})$ and any permutation σ , we have from (6.9)

$$\rho(\sigma(\mathbf{j}), \sigma(\boldsymbol{\pi})) = \rho(\mathbf{j}, \boldsymbol{\pi})$$

i.e., ρ is invariant to permutations. Hence, we can rewrite (6.21)–(6.11) as

$$\det(-\mathbf{J}(p)) = \sum_{(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0} D(\mathbf{j}, \boldsymbol{\pi}) \rho(\mathbf{j}, \boldsymbol{\pi}) \quad (6.13)$$

where

$$D(\mathbf{j}, \boldsymbol{\pi}) = \sum_{\boldsymbol{\sigma}} \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l})) \operatorname{sgn} \mathbf{k} \mu(\boldsymbol{\sigma}(\mathbf{j})) \quad (6.14)$$

In the above, L -vectors $\boldsymbol{\sigma}$ and \mathbf{k} are permutations.

The next lemma converts a condition on $\boldsymbol{\sigma}(\boldsymbol{\pi})$ into one on $\boldsymbol{\pi}$. It follows directly from the definition of permutation.

Lemma 6.2. *For any $\boldsymbol{\pi}$ and any permutations $\boldsymbol{\sigma}, \mathbf{k}$, we have*

$$\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l}) \Leftrightarrow \boldsymbol{\pi} \in \Pi(\boldsymbol{\sigma}^{-1} \mathbf{k}, \boldsymbol{\sigma}^{-1} \mathbf{l})$$

i.e., $[\boldsymbol{\sigma}(\boldsymbol{\pi})]_l \in \Pi(k_l, l)$ for all l if and only if $\pi^l \in \Pi(k_{\sigma_l^{-1}}, \sigma_l^{-1})$ for all l .

¹The one-one correspondence fails to hold for permutations not in Θ .

Applying Lemma 6.2 to (6.26), we have

$$D(\mathbf{j}, \boldsymbol{\pi}) = \sum_{\boldsymbol{\sigma}} \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\pi} \in \Pi(\boldsymbol{\sigma}^{-1}\mathbf{k}, \boldsymbol{\sigma}^{-1})) \operatorname{sgn}\mathbf{k} \mu(\boldsymbol{\sigma}(\mathbf{j}))$$

Since \mathbf{k} , and hence $\boldsymbol{\sigma}^{-1}\mathbf{k}$, range over all possible permutations, we can replace the index variable $\boldsymbol{\sigma}^{-1}\mathbf{k}$ by \mathbf{k} to get

$$D(\mathbf{j}, \boldsymbol{\pi}) = \sum_{\boldsymbol{\sigma}} \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \boldsymbol{\sigma}^{-1})) \operatorname{sgn}(\mathbf{k}\boldsymbol{\sigma}) \mu(\boldsymbol{\sigma}(\mathbf{j})) \quad (6.15)$$

We now use (6.15) to derive a sufficient condition under which $D(\mathbf{j}, \boldsymbol{\pi})$ are nonnegative for all permutationally distinct $(\mathbf{j}, \boldsymbol{\pi})$. The main idea is to show that for every negative term in the summation in (6.15), either it can be exactly cancelled by a positive term, or we can find two positive terms whose sum has a larger or equal magnitude under the given condition. This lemma directly implies Theorem 2.8.

Lemma 6.3. *Suppose assumptions A1–A3 hold. Suppose for any $\mathbf{j} \in \{1, \dots, J\}^L$ and any permutations $\boldsymbol{\sigma}, \mathbf{k}, \mathbf{n}$ in $\{1, \dots, L\}^L$, we have for a regular network*

$$\mu(\mathbf{k}(\mathbf{j})) + \mu(\mathbf{n}(\mathbf{j})) \geq \mu(\boldsymbol{\sigma}(\mathbf{j}))$$

Then, for all $(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0$, $D(\mathbf{j}, \boldsymbol{\pi}) \geq 0$.

Proof. Fix any $(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0$. Each term in (6.15) is indexed by a pair $(\boldsymbol{\sigma}, \mathbf{k})$.

Fix also a permutation $\boldsymbol{\sigma}$ in (6.15). Suppose there is only one permutation \mathbf{k} for which the term indexed by $(\boldsymbol{\sigma}, \mathbf{k})$ has a negative sign given by $\mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \boldsymbol{\sigma}^{-1})) \operatorname{sgn}(\mathbf{k}\boldsymbol{\sigma}) = -1$. This term is then $-\mu(\boldsymbol{\sigma}(\mathbf{j})) < 0$. Since the summation over \mathbf{k} ranges over all permutations, we can find a positive term, indexed by $(\boldsymbol{\sigma}, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \boldsymbol{\sigma}^{-1}$, that exactly cancels this negative term. This is because $\mathbf{1}(\boldsymbol{\pi} \in \Pi(\hat{\mathbf{k}}, \boldsymbol{\sigma}^{-1}))$ is always 1 and $\operatorname{sgn}(\hat{\mathbf{k}}\boldsymbol{\sigma}) = \operatorname{sgn}\mathbf{l} = 1$, yielding the term $\mu(\boldsymbol{\sigma}(\mathbf{j}))$. Hence we have shown that, given $\boldsymbol{\sigma}$, if there is only one \mathbf{k} that yields a negative term, then it is always cancelled by another positive term indexed by $(\boldsymbol{\sigma}, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \boldsymbol{\sigma}^{-1}$.

Given a σ , suppose now there are two permutations \mathbf{k}, \mathbf{n} for which

$$\pi \in \Pi(\mathbf{k}, \sigma^{-1}) \quad \text{and} \quad \pi \in \Pi(\mathbf{n}, \sigma^{-1}) \quad (6.16)$$

and $\text{sgn}(\mathbf{k}\sigma) = \text{sgn}(\mathbf{n}\sigma) = -1$. Each of (σ, \mathbf{k}) and (σ, \mathbf{n}) yields a negative term $-\mu(\sigma(\mathbf{j}))$ in the summation in (6.15). Notice that (6.27) says that, for all $l = 1, \dots, L$, the path π^l contains link pairs (k_l, σ_l^{-1}) and (n_l, σ_l^{-1}) . Hence π^l also pass through link pairs $(\sigma_l^{-1}, \sigma_l^{-1})$, (k_l, n_l) and (n_l, k_l) , i.e.,

$$\pi \in \Pi(\sigma^{-1}, \sigma^{-1}) \quad (6.17)$$

$$\pi \in \Pi(\mathbf{k}, \mathbf{n}), \quad \pi \in \Pi(\mathbf{n}, \mathbf{k}) \quad (6.18)$$

(6.29) implies that there is a positive term in the summation in (6.15) indexed by $(\sigma, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \sigma^{-1}$:

$$\text{sgn}(\sigma^{-1}\sigma)\mu(\sigma(\mathbf{j})) = \mu(\sigma(\mathbf{j})) > 0$$

It cancels the negative term $-\mu(\sigma(\mathbf{j}))$ in the summation indexed by (σ, \mathbf{k}) .

To deal with the negative term $-\mu(\sigma(\mathbf{j}))$ indexed by (σ, \mathbf{n}) , note that (6.30) implies that there are two nonzero terms in the summation, indexed by $(\mathbf{n}^{-1}, \mathbf{k})$ and $(\mathbf{k}^{-1}, \mathbf{n})$, that we now argue are positive. Indeed the term indexed by $(\mathbf{n}^{-1}, \mathbf{k})$ is

$$\begin{aligned} \text{sgn}(\mathbf{k}\mathbf{n}^{-1}) \mu(\mathbf{n}^{-1}(\mathbf{j})) &= \text{sgn}(\mathbf{k}\sigma(\mathbf{n}\sigma)^{-1}) \mu(\mathbf{n}^{-1}(\mathbf{j})) \\ &= \text{sgn}(\mathbf{k}\sigma) \text{sgn}(\mathbf{n}\sigma)^{-1} \mu(\mathbf{n}^{-1}(\mathbf{j})) \\ &= \mu(\mathbf{n}^{-1}(\mathbf{j})) > 0 \end{aligned}$$

where we have used the hypothesis that $\text{sgn}(\mathbf{k}\sigma) = -1$ and $\text{sgn}(\mathbf{n}\sigma)^{-1} = \text{sgn}(\mathbf{n}\sigma) = -1$. Similarly, the term with index $(\mathbf{k}^{-1}, \mathbf{n})$ is $\mu(\mathbf{k}^{-1}(\mathbf{j}))$. The hypothesis of the lemma implies that

$$\mu(\mathbf{n}^{-1}(\mathbf{j})) + \mu(\mathbf{k}^{-1}(\mathbf{j})) - \mu(\sigma(\mathbf{j})) \geq 0$$

Hence, we have shown that, given σ , if there are two negative terms in the summation in (6.15) indexed by (σ, \mathbf{k}) and (σ, \mathbf{n}) , then we can always find three positive terms, indexed by, (σ, σ^{-1}) , $(\mathbf{n}^{-1}, \mathbf{k})$ and $(\mathbf{k}^{-1}, \mathbf{n})$, so that the sum of these five terms is nonnegative.

If there are more than two negative terms, take any *additional* negative term, indexed by, say, $(\sigma, \hat{\mathbf{n}})$. The same argument shows that it will be compensated by the two (unique) positive terms indexed by $(\hat{\mathbf{n}}^{-1}, \mathbf{k})$ and $(\mathbf{k}^{-1}, \hat{\mathbf{n}})$. This completes the proof. \square

Since the network is regular, $\det(-\mathbf{J}(p)) \neq 0$. Lemma 6.3, together with (6.13), implies that $\det(-\mathbf{J}(p)) > 0$, or equivalently, $I(p) = (-1)^L$ for any $p \in E$, under the condition of the lemma. Theorem 2.8 then follows from Corollary 2.7. An illustration for the proof of Lemma 6.3 via a concrete example ($L = 3, J = 2$) can be found in Appendix 6.3.

Remark: The sufficient condition in Theorem 2.8 can be conservative because many r_π^j may be zero (no source of type j takes path π).

6.3 Proof of Theorem 2.10

Proof. It is straightforward to check that only the following six $\rho(\mathbf{j}, \boldsymbol{\pi})$ in (6.13) can have negative coefficients $D(\mathbf{j}, \boldsymbol{\pi})$:

$$\begin{aligned}
& (\lambda_2 + \lambda_3 - \lambda_1) \dot{m}_1^2 \dot{m}_2^2 \dot{m}_3^2 r_{111}^1 r_{101}^2 r_{110}^2 & (6.19) \\
& (\lambda_1 + \lambda_3 - \lambda_2) \dot{m}_1^2 \dot{m}_2^2 \dot{m}_3^2 r_{111}^1 r_{011}^2 r_{110}^2 \\
& (\lambda_1 + \lambda_2 - \lambda_3) \dot{m}_1^2 \dot{m}_2^2 \dot{m}_3^2 r_{111}^1 r_{011}^2 r_{101}^2 \\
& \left(\frac{1}{\lambda_2} + \frac{1}{\lambda_3} - \frac{1}{\lambda_1}\right) \dot{m}_1^1 \dot{m}_2^1 \dot{m}_3^1 r_{111}^2 r_{101}^1 r_{110}^1 \\
& \left(\frac{1}{\lambda_1} + \frac{1}{\lambda_3} - \frac{1}{\lambda_2}\right) \dot{m}_1^1 \dot{m}_2^1 \dot{m}_3^1 r_{111}^2 r_{011}^1 r_{110}^1 \\
& \left(\frac{1}{\lambda_1} + \frac{1}{\lambda_2} - \frac{1}{\lambda_3}\right) \dot{m}_1^1 \dot{m}_2^1 \dot{m}_3^1 r_{111}^2 r_{011}^1 r_{101}^1
\end{aligned}$$

The condition in the theorem guarantees that these terms are all nonnegative. By (6.13), $\det(-\mathbf{J}(p)) \geq 0$. Since the network is regular, we have $\det(-\mathbf{J}(p)) > 0$ for all equilibria p . Hence the equilibrium is globally unique. \square

We close this section by illustrating how we determine the coefficient $D(\mathbf{j}, \boldsymbol{\pi})$ in the proof of Lemma 6.3. Consider the term for $\rho(\mathbf{j}, \boldsymbol{\pi}) = r_{111}^1 r_{101}^2 r_{110}^2$ in (6.19). Here $\mathbf{j} = (1, 2, 2)$ and $\boldsymbol{\pi} = ((111), (101), (110))$. By (6.26), we need to look at the sum over $\boldsymbol{\sigma}$ and \mathbf{k} . First, look at $\boldsymbol{\sigma} = (3, 1, 2)$, the only \mathbf{k} such that $\mathbf{1}(\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l})) = 1$ and $\text{sgn}\mathbf{k} = -1$ is $\mathbf{k} = (2, 1, 3)$. By the argument in the proof of Lemma 6.3, if we let $\mathbf{k} = \mathbf{l} = (1, 2, 3)$, we have $\mathbf{1}(\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l})) = 1$ and $\text{sgn}\mathbf{k} = 1$ and the sum of these two terms in (6.26) is zero.

We can visualize this operation as follows: Each entry of $-\mathbf{J}(p)$ is a sum of $\dot{m}_i^j r_\pi^j$ with appropriate signs. When we expand its determinant, we obtain, from (6.13)–(6.26), a sum, over a set of source types \mathbf{j} , paths $\boldsymbol{\pi}$ and permutations $\boldsymbol{\sigma}, \mathbf{k}$, of terms $\rho(\mathbf{j}, \boldsymbol{\pi})$ that are products of r_π^j . Hence we can identify each term in (6.13)–(6.26), indexed by $(\mathbf{j}, \boldsymbol{\pi}, \boldsymbol{\sigma}, \mathbf{k})$, with the original position in $-\mathbf{J}(p)$ of each constituent r_π^j in $\rho(\mathbf{j}, \boldsymbol{\pi})$. This is illustrated below: The negative term

$$\begin{bmatrix} & r_{111}^1 & \\ r_{110}^2 & & \\ & & r_{101}^2 \end{bmatrix} (\boldsymbol{\sigma} = (3, 1, 2), \mathbf{k} = (2, 1, 3), \text{sgn}\mathbf{k} = -1)$$

is cancelled exactly by the positive term

$$\begin{bmatrix} r_{110}^2 & & \\ & r_{111}^1 & \\ & & r_{101}^2 \end{bmatrix} (\boldsymbol{\sigma} = (3, 1, 2), \mathbf{k} = (1, 2, 3), \text{sgn}\mathbf{k} = 1)$$

Similarly, we have the following two terms that cancel one another:

$$\begin{bmatrix} & & r_{111}^1 \\ & r_{110}^2 & \\ r_{101}^2 & & \end{bmatrix} (\boldsymbol{\sigma} = (2, 3, 1), \mathbf{k} = (3, 2, 1), \text{sgn}\mathbf{k} = -1)$$

$$\begin{bmatrix} r_{101}^2 & & \\ & r_{110}^2 & \\ & & r_{111}^1 \end{bmatrix} (\boldsymbol{\sigma} = (2, 3, 1), \mathbf{k} = (1, 2, 3), \text{sgn}\mathbf{k} = 1)$$

Now consider $\boldsymbol{\sigma} = (1, 3, 2)$. We have the following two terms with $\text{sgn}\mathbf{k} = -1$.

$$\begin{bmatrix} & & r_{101}^2 \\ & r_{110}^2 & \\ r_{111}^1 & & \end{bmatrix} (\boldsymbol{\sigma} = (1, 3, 2), \mathbf{k} = (3, 2, 1), \text{sgn}\mathbf{k} = -1)$$

$$\begin{bmatrix} & r_{110}^2 & \\ r_{111}^1 & & \\ & & r_{101}^2 \end{bmatrix} (\boldsymbol{\sigma} = (1, 3, 2), \mathbf{k} = (2, 1, 3), \text{sgn}\mathbf{k} = -1)$$

Setting $\mathbf{k} = \mathbf{l} = (1, 2, 3)$ gives the following positive term:

$$\begin{bmatrix} r_{111}^1 & & \\ & r_{110}^2 & \\ & & r_{101}^2 \end{bmatrix} (\boldsymbol{\sigma} = (1, 3, 2), \mathbf{k} = (1, 2, 3), \text{sgn}\mathbf{k} = 1)$$

As described in the proof of Lemma 6.3, we can find two positive terms indexed by some $(\boldsymbol{\sigma}, \mathbf{k})$. One is $(\boldsymbol{\sigma} = (3, 2, 1)(1, 3, 2) = (2, 3, 1), \mathbf{k} = (3, 2, 1)^{-1}(2, 1, 3) = (3, 1, 2))$ and the other is $(\boldsymbol{\sigma} = (2, 1, 3)(1, 3, 2) = (3, 1, 2), \mathbf{k} = (2, 1, 3)^{-1}(3, 2, 1) = (2, 3, 1))$. They can be visualized as the following:

$$\begin{bmatrix} & r_{110}^2 & \\ & & r_{111}^1 \\ r_{101}^2 & & \end{bmatrix} (\boldsymbol{\sigma} = (2, 3, 1), \mathbf{k} = (3, 1, 2), \text{sgn}\mathbf{k} = 1)$$

$$\begin{bmatrix} & & r_{101}^2 \\ r_{110}^2 & & \\ & r_{111}^1 & \end{bmatrix} (\boldsymbol{\sigma} = (3, 1, 2), \mathbf{k} = (2, 3, 1), \text{sgn}\mathbf{k} = 1)$$

For $\boldsymbol{\pi} = ((111), (101), (110))$, we can actually verify that only the nine terms dis-

cussed above have $\mathbf{1}(\sigma(\pi) \in \Pi(\mathbf{k}, \mathbf{l})) \text{sgn} \mathbf{k} \neq 0$. Therefore, the coefficient $D(\mathbf{j}, \pi) = \mu((2, 3, 1)(\mathbf{j})) + \mu((3, 1, 2)(\mathbf{j})) - \mu((1, 3, 2)(\mathbf{j}))$. Noting $\mathbf{j} = (1, 2, 2)$, we finally get

$$\begin{aligned} D(\mathbf{j}, \pi) &= \mu((2, 2, 1)) + \mu((2, 1, 2)) - \mu((1, 2, 2)) \\ &= \dot{m}_1^2 \dot{m}_2^2 \dot{m}_3^1 + \dot{m}_1^1 \dot{m}_2^1 \dot{m}_3^2 - \dot{m}_1^1 \dot{m}_2^2 \dot{m}_3^2 \\ &= (\lambda_3 + \lambda_2 - \lambda_1) \dot{m}_1^2 \dot{m}_2^2 \dot{m}_3^2 \end{aligned}$$

6.4 Smallest Network with Multiple Equilibria

Example 6.1: a two-link network with non-unique equilibria

In this example, we assume that all of the sources use the same utility function defined as

$$U_i^j(x_i^j) = -\frac{1}{2}(1 - x_i^j)^2$$

The network has two links with capacity vector $c = [1, 1]$. The corresponding routing matrices for these two protocols are

$$R^1 = R^2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

We use linear price mapping functions $m^j(p) = K^j p$, $j = 1, 2$, where K^j are 2×2 matrices given by

$$K^1 = I, \quad K^2 = \text{diag}(1, 3)$$

As for Example 1, we check the matrix

$$\sum_{i=1}^2 R^i (R^i)^T K^i = \begin{bmatrix} 1 & 4 \\ 1 & 4 \end{bmatrix}$$

which has determinant 0, implying multiple equilibria. It is easy to verify that the following

points are all equilibria:

$$p_1 = \epsilon, \quad p_2 = 1/4 - \epsilon/2, \quad \text{where } \epsilon \in [0, 1/2]$$

The corresponding rates are

$$x_1^1 = 3/4 - \epsilon/2, \quad x_1^2 = 1/4 + \epsilon/2$$

The capacity constraints are all tight. □

Remarks: Note that even with a single protocol, the example above has non-unique equilibrium price vectors since the routing matrix is not full rank. However, in that case, the equilibrium rate vector is unique, unlike the case of multiple protocols.

6.5 Proof of Lemma 2.16

Proof. Suppose that λ is an eigenvalue of $B + \mathbf{1}\gamma^T$, then $\text{diag}(\beta_i - \lambda) + \mathbf{1}\gamma^T$ is singular. If $\lambda = \beta_i$ for certain i , then, since $\beta_i > 0$, λ is positive. Otherwise the following matrix is also singular:

$$I + \text{diag}\left(\frac{1}{\beta_i - \lambda}\right) \mathbf{1}\gamma^T \tag{6.20}$$

The rank of matrix $\text{diag}(1/(\beta_i - \lambda))\mathbf{1}\gamma^T$ is 1. Moreover it has only one nonzero eigenvalue equal to $\sum_i \gamma_i/(\beta_i - \lambda)$. For the matrix in (6.20) to be singular, it must have a zero eigenvalue, and this is possible if and only if

$$\sum_i \frac{\gamma_i}{\beta_i - \lambda_i} = -1$$

The real part of $\gamma_i/(\beta_i - \lambda_i)$ is $\gamma_i(\beta_i - \text{Re}\lambda)/|\beta_i - \lambda|^2$. If $\text{Re}\lambda \leq 0$, the sum of the real part of $\gamma_i/(\beta_i - \lambda_i)$ cannot be -1 . So we must have $\text{Re}\lambda > 0$. □

6.6 Proof of Theorem 3.8

For a real matrix A , if all its principle minors are positive, A is called a P -matrix. If $a_{ii} \geq 0$, $a_{ij} \leq 0$, then A is called an M -matrix.² Clearly if a P -matrix is symmetric, then it is positive definite and hence stable. However, the Jacobian matrix in our problem is not symmetric due to the fact that multiple protocols exist, which is the main difficulty of setting up stability. Before getting into the main proof, we state three lemmas here. One is referred to [10] for other related results.

Lemma 6.4. *If A is a P -matrix and also an M -matrix, then all its eigenvalues have positive real parts.*

Proof: Choose $\alpha \geq \max_i a_{ii}$, then $G = \alpha I - A \geq 0$. By the classical Perron-Frobenius theorem [10], we can find a non-negative γ that is an eigenvalue of G with the largest modulus. Let $\mu = \alpha - \gamma$, then μ is an eigenvalue of A . Let λ be another eigenvalue of A and $\lambda \neq \mu$. Then $\beta = \alpha - \lambda$ is an eigenvalue of G . Hence $|\beta| \leq \gamma$. As $\beta \neq \gamma$, we have $\Re(\beta) < \gamma$ where $\Re(\beta)$ is the real part of β . Then $\Re\lambda = \alpha - \Re(\beta) > \alpha - \gamma = \mu$. In short, A has an eigenvalue μ , which is real and whose real part is less than that of any other different eigenvalues of A .

On the other hand, as A is also a P -matrix, it is well known that A cannot have a nonpositive real eigenvalues, which can be clearly seen from the characteristic polynomial of A . Now suppose there is an eigenvalue λ of A whose real part is not positive, as we showed in the last paragraph, there is another eigenvalue μ of A , which is real and $\mu < \Re(\lambda) \leq 0$. This contradicts the fact that A is a P -matrix and hence all eigenvalues of A have positive real parts.

□

Remark: It is well known in mathematical economics that if the Jacobian matrix is an M -matrix then it is also a P -matrix. The main reason is there is another property of the

²We use terminologies from the mathematics community. As closely related problems are extensively studied in economics, there are parallel terminologies. e.g., if A is a P -matrix, then $-A$ is called Hecksian; an M -matrix with positive diagonals is referred to as Metzlerian or as having gross substitution property.

Jacobian matrix that can be obtained by Walras's law or the fact that the demand function is of homogeneity zero. In our problem, although complementary slackness condition can generate the same property as Walras's law does (only at equilibrium), our Jacobian matrix is not an M -matrix. The M -matrix on which we will use this lemma does not have that additional property and hence needs an additional assumption to be a P -matrix here. Checking whether a matrix is a P -matrix or not is known to be NP-complete [16]. However, using the structure of our problem we can still analytically check the P -matrix condition and hence set up our main result in the next section.

Let e be the column vector $e = [1, 1, \dots, 1]^T$.

Lemma 6.5. *If A is an M -matrix and all its eigenvalues have positive real parts, then there is an $D = \text{diag}[d_1, \dots, d_n]$, $d_i > 0$ for all i , such that $D^{-1}ADe = h > 0$. In other words, A is diagonally dominant.*

Proof: As all eigenvalues of A have positive real parts, then the real one with the smallest real part is greater than 0, i.e., $\mu > 0$. Still define $G = \alpha I - A$ with $\alpha \geq \max_i a_{ii}$ and let γ be an eigenvalue of G with the largest modulus. We then have $\alpha = \mu + \gamma > \gamma$. Again by the Perron-Frobenius theorem, $(\alpha I - G)^{-1} = A^{-1} \geq 0$.

Pick any positive vector $v > 0$, we can conclude $A^{-1}v > 0$ as every row of the nonsingular matrix A^{-1} contains at least one positive element. Define a diagonal matrix $D > 0$ uniquely by $De = A^{-1}v$. Then $D^{-1}ADe = D^{-1}v > 0$. \square

Remark: The above claim provides a sufficient condition to test whether an M -matrix is diagonally dominant, the condition is actually also necessary [66].

For a matrix A , we define its comparison matrix $M(A) = (m_{ij})$ by setting $m_{ii} = |a_{ii}|$, and $m_{ij} = -|a_{ij}|$ if $i \neq j$. Clearly $M(A)$ is an M -matrix. The following lemma points out a simple yet important fact that relates diagonal dominance property of A with positive diagonal entries and that of $M(A)$.

Lemma 6.6. *Suppose all diagonal entries of A are positive. If there is an $D = \text{diag}[d_1, \dots, d_n]$, $d_i > 0$ for all i , such that $D^{-1}M(A)De = h > 0$. Then $D^{-1}ADe > 0$, i.e., A is also diagonally dominant.*

Proof: For any i , we have

$$(D^{-1}ADe)_i = a_{ii} - \sum_{j \neq i} \frac{d_j}{d_i} |a_{ij}| = (D^{-1}M(A)De)_i > 0$$

□

We now state the proof of Theorem 3.8.

Proof: We need to show all eigenvalues of $-\mathbf{J}$ have positive real parts. It is enough to show $-\mathbf{J}$ is diagonally dominant and by Lemma 6.6 we only need to show $M(-\mathbf{J})$ is diagonally dominant as all diagonal entries of $-\mathbf{J}$ are positive (each link has at least one flow using it). Using Lemma 6.5, it suffices to show that $M(-\mathbf{J})$ is positive stable, which then can be reduced to check whether $M(-\mathbf{J})$ is a P -matrix by Lemma 6.4. By similar arguments in [75], it is enough to show $\det(M(-\mathbf{J})) > 0$, which will be done in the remainder of the proof. Following section 6.2, we have

$$\det(-\mathbf{J}(p)) = \sum_{\mathbf{j}} \sum_{\boldsymbol{\pi}} C(\mathbf{j}, \boldsymbol{\pi}) \rho(\mathbf{j}, \boldsymbol{\pi}) \quad (6.21)$$

where

$$C(\mathbf{j}, \boldsymbol{\pi}) := \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \mathbf{l})) \operatorname{sgn} \mathbf{k} \mu(\mathbf{j}) \quad (6.22)$$

For any permutation \mathbf{k} , Define $L_{\mathbf{k}}^+ = \{l | k_l = l\}$ and $L_{\mathbf{k}}^- = \{l | k_l \neq l\}$. We then have

$$\det(M(-\mathbf{J})) = \sum_{\mathbf{j}} \sum_{\boldsymbol{\pi}} G(\mathbf{j}, \boldsymbol{\pi}) \rho(\mathbf{j}, \boldsymbol{\pi}) \quad (6.23)$$

where

$$G(\mathbf{j}, \boldsymbol{\pi}) := \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\pi} \in \Pi(\mathbf{k}, \mathbf{l})) \operatorname{sgn} \mathbf{k} (-1)^{|L_{\mathbf{k}}^-|} \mu(\mathbf{j}) \quad (6.24)$$

Hence $\det M(-\mathbf{J}(p))$ is a summation, over the index $(\mathbf{j}, \boldsymbol{\pi})$, of terms $\rho(\mathbf{j}, \boldsymbol{\pi})$ with coeffi-

icients $G(\mathbf{j}, \boldsymbol{\pi})$.

Then let Θ_0 be the largest subset of the set of all possible $(\mathbf{j}, \boldsymbol{\pi})$'s that is *permutationally distinct*, i.e., no vector in Θ_0 is a permutation of another vector in Θ_0 . We then have

$$\det(M(-\mathbf{J}(p))) = \sum_{(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0} H(\mathbf{j}, \boldsymbol{\pi}) \rho(\mathbf{j}, \boldsymbol{\pi}) \quad (6.25)$$

$$H(\mathbf{j}, \boldsymbol{\pi}) = \sum_{\boldsymbol{\sigma} \in \Sigma(\mathbf{j}, \boldsymbol{\pi})} \sum_{\mathbf{k}} \mathbf{1}(\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l})) T \quad (6.26)$$

where

$$T = \text{sgn} \mathbf{k} (-1)^{|\mathbf{k}^-|} \mu(\boldsymbol{\sigma}(\mathbf{j}))$$

and $\Sigma(\mathbf{j}, \boldsymbol{\pi})$ is the largest subset of the set of all permutations $\boldsymbol{\sigma}$ that generate distinct $\boldsymbol{\sigma}(\mathbf{j}, \boldsymbol{\pi})$.

We now use (6.26) to derive a sufficient condition under which $H(\mathbf{j}, \boldsymbol{\pi})$ are nonnegative for all permutationally distinct $(\mathbf{j}, \boldsymbol{\pi})$. The main idea is to show that for every negative term in the summation in (6.26), either it can be exactly cancelled by a positive term, or we can find two positive terms whose sum has a larger or equal magnitude under the given condition. This lemma directly implies Theorem 3.8.

Lemma 6.7. *Suppose assumptions A1–A3 hold. Suppose for any $\mathbf{j} \in \{1, \dots, J\}^L$ and any permutations $\boldsymbol{\sigma}, \mathbf{k}, \mathbf{n}$ in $\{1, \dots, L\}^L$, we have for a regular network*

$$\mu(\mathbf{k}(\mathbf{j})) + \mu(\mathbf{n}(\mathbf{j})) \geq \mu(\boldsymbol{\sigma}(\mathbf{j}))$$

Then, for all $(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0$, $H(\mathbf{j}, \boldsymbol{\pi}) \geq 0$.

Proof. Fix any $(\mathbf{j}, \boldsymbol{\pi}) \in \Theta_0$. Each term in (6.26) is indexed by a pair $(\boldsymbol{\sigma}, \mathbf{k})$. Fix also a permutation $\boldsymbol{\sigma}$ in (6.26). Suppose there is only one permutation \mathbf{k} for which the term indexed by $(\boldsymbol{\sigma}, \mathbf{k})$ has a negative sign given by $\mathbf{1}(\boldsymbol{\sigma}(\boldsymbol{\pi}) \in \Pi(\mathbf{k}, \mathbf{l})) \text{sgn}(\mathbf{k}) (-1)^{|\mathbf{k}^-|} = -1$. This term is then $-\mu(\boldsymbol{\sigma}(\mathbf{j})) < 0$. Since the summation over \mathbf{k} ranges over all permutations,

we can find a positive term, indexed by $(\sigma, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \mathbf{l}$, that exactly cancels this negative term. This is because $\mathbf{1}(\sigma(\pi) \in \Pi(\mathbf{l}, \mathbf{l}))$ is always 1 and $\text{sgn}(\mathbf{l})(-1)^{|L_{\mathbf{l}}^-|} = 1$, yielding the term $\mu(\sigma(\mathbf{j}))$. Hence we have shown that, given σ , if there is only one \mathbf{k} that yields a negative term, then it is always cancelled by another positive term indexed by $(\sigma, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \mathbf{l}$.

Given a σ , suppose now there are two permutations \mathbf{k}, \mathbf{n} for which

$$\sigma(\pi) \in \Pi(\mathbf{k}, \mathbf{l}) \quad \text{and} \quad \sigma(\pi) \in \Pi(\mathbf{n}, \mathbf{l}) \quad (6.27)$$

and

$$\text{sgn}(\mathbf{k})(-1)^{|L_{\mathbf{k}}^-|} = \text{sgn}(\mathbf{n})(-1)^{|L_{\mathbf{n}}^-|} = -1 \quad (6.28)$$

Each of (σ, \mathbf{k}) and (σ, \mathbf{n}) yields a negative term $-\mu(\sigma(\mathbf{j}))$ in the summation in (6.26). Notice that (6.27) says that, for all $l = 1, \dots, L$, paths $\sigma(\pi)^l$ contains link pairs (k_l, l) and (n_l, l) . Hence $\sigma(\pi)^l$ also pass through link pairs (l, l) , (k_l, n_l) and (n_l, k_l) , i.e.,

$$\sigma(\pi) \in \Pi(\mathbf{l}, \mathbf{l}) \quad (6.29)$$

$$\sigma(\pi) \in \Pi(\mathbf{k}, \mathbf{n}), \quad \sigma(\pi) \in \Pi(\mathbf{n}, \mathbf{k}) \quad (6.30)$$

(6.29) implies that there is a positive term in the summation in (6.26) indexed by $(\sigma, \hat{\mathbf{k}})$ with $\hat{\mathbf{k}} = \mathbf{l}$:

$$\text{sgn}(\mathbf{l})(-1)^{|L_{\mathbf{l}}^-|} \mu(\sigma(\mathbf{j})) = \mu(\sigma(\mathbf{j})) > 0$$

It cancels the negative term $-\mu(\sigma(\mathbf{j}))$ in the summation indexed by (σ, \mathbf{k}) .

To deal with the negative term $-\mu(\sigma(\mathbf{j}))$ indexed by (σ, \mathbf{n}) , note that (6.30) implies that there are two nonzero terms in the summation, indexed by $(\mathbf{n}^{-1}\sigma, \mathbf{n}^{-1}\mathbf{k})$ and $(\mathbf{k}^{-1}\sigma, \mathbf{k}^{-1}\mathbf{n})$, that we now argue are positive. Indeed the term indexed by $(\mathbf{n}^{-1}\sigma, \mathbf{n}^{-1}\mathbf{k})$

is $\text{sgn}(\mathbf{n}^{-1}\mathbf{k}) (-1)^{|L_{\mathbf{n}^{-1}\mathbf{k}}^-|} \mu(\mathbf{n}^{-1}(\mathbf{j}))$. We further have

$$\begin{aligned} |L_{\mathbf{n}^{-1}\mathbf{k}}^-| &= |L_{\mathbf{k}}^- \cup L_{\mathbf{n}}^-| - |(L_{\mathbf{k}}^- \cap L_{\mathbf{n}}^-)| \\ &= |L_{\mathbf{k}}^-| + |L_{\mathbf{n}}^-| - 2|(L_{\mathbf{k}}^- \cap L_{\mathbf{n}}^-)| \end{aligned} \quad (6.31)$$

Hence

$$\begin{aligned} \text{sgn}(\mathbf{n}^{-1}\mathbf{k}) (-1)^{|L_{\mathbf{n}^{-1}\mathbf{k}}^-|} &= \text{sgn}(\mathbf{n})\text{sgn}(\mathbf{k}) (-1)^{|L_{\mathbf{k}}^-|} (-1)^{|L_{\mathbf{n}}^-|} \\ &= 1 \end{aligned}$$

The last equality follows from (6.28). Therefore,

$$\text{sgn}(\mathbf{n}^{-1}\mathbf{k}) (-1)^{|L_{\mathbf{n}^{-1}\mathbf{k}}^-|} = \mu(\mathbf{n}^{-1}(\mathbf{j})) > 0$$

Similarly, the term with index $(\mathbf{k}^{-1}\boldsymbol{\sigma}, \mathbf{k}^{-1}\mathbf{n})$ is $\mu(\mathbf{k}^{-1}(\mathbf{j}))$. The hypothesis of the lemma implies that

$$\mu(\mathbf{n}^{-1}(\mathbf{j})) + \mu(\mathbf{k}^{-1}(\mathbf{j})) - \mu(\boldsymbol{\sigma}(\mathbf{j})) \geq 0$$

Hence, given $\boldsymbol{\sigma}$, if there are two negative terms in the summation in (6.26) indexed by $(\boldsymbol{\sigma}, \mathbf{k})$ and $(\boldsymbol{\sigma}, \mathbf{n})$, then we can always find three positive terms, indexed by $(\boldsymbol{\sigma}, \mathbf{l})$, $(\mathbf{n}^{-1}\boldsymbol{\sigma}, \mathbf{n}^{-1}\mathbf{k})$ and $(\mathbf{k}^{-1}\boldsymbol{\sigma}, \mathbf{k}^{-1}\mathbf{n})$, so that the sum of these five terms are nonnegative.

If there are more than two negative terms, take any *additional* negative term, indexed by, say, $(\boldsymbol{\sigma}, \hat{\mathbf{n}})$. The same argument shows that it will be compensated by the two (unique) positive terms indexed by $(\hat{\mathbf{n}}^{-1}\boldsymbol{\sigma}, \hat{\mathbf{n}}^{-1}\mathbf{k})$ and $(\mathbf{k}^{-1}\boldsymbol{\sigma}, \mathbf{k}^{-1}\hat{\mathbf{n}})$. This completes the proof. \square

6.7 Proof of Theorem 4.2

Proof. The analysis uses the standard singular perturbation stability results [40]. By plugging

$$\mu_i^j = \frac{q_i^j}{\sum_{l \in L(i,j)} p_l}$$

into (4.9), we get the reduced system

$$\dot{p}_l(t) = y_l(p(t)) - c_l \quad (6.32)$$

where

$$x_i^j = (U_i^j)^{t-1} \left(\sum_{l \in L(i,j)} p_l \right)$$

In other words, the reduced system has exactly the same dynamics as a network with a single congestion price. Hence it is globally asymptotically stable, which can be verified by choosing a Lyapunov function $V_1(p) = D(p) - \min(D(p))$, where $D(p)$ is the dual function. Let $\phi_1(p) = |(y(p) - c)|$, which is a positive definite function. Then

$$\frac{\partial V_1}{\partial p}(\dot{p}) = -(y(p) - c)^2 \leq -\phi_1^2(p) \quad (6.33)$$

Let

$$z_i^j = \mu_i^j - \frac{q_i^j}{\sum_{l \in L(i,j)} p_l}$$

and scale time t by $\tau = \frac{t}{\epsilon}$, we get the boundary-layer system

$$\frac{dz_i^j}{d\tau} = -z_i^j \quad (6.34)$$

By using a Lyapunov function $V_2(w) = \frac{1}{2} w^T * w$, one can show the boundary-layer system

is globally asymptotically stable. Let $\phi_2(z) = \text{sqrt}(z^T z)$, which is a positive definite function. Then

$$\frac{\partial V_2}{\partial z} \frac{dz_i^j}{d\tau} = -z^T z \leq -\phi_2^2(z) \quad (6.35)$$

Following arguments in [40], to prove the existence of $\epsilon^* > 0$ in the theorem, we only need to further check

$$\left\| \frac{\partial V_2}{\partial z} \right\| \leq k_1 \phi_2(z)$$

and

$$\|y(p, z) - y(p)\| \leq k_2 \phi_2(z)$$

where k_1 and k_2 are constants. The first one is true by the definition of V_2 and we can simply choose $k_1 = 1$, the second requirement is also met as y is Lipschitz continuous [46]. □

Bibliography

- [1] Linux kernel document file: Documentation/networking/tcp.txt.
- [2] WAN in Lab. <http://wil.cs.caltech.edu>.
- [3] D. Acemoglu, A. Ozdaglar, A. Simsek, and A. Tang. On the stability of P-matrices. *working paper*, 2006.
- [4] F. Ackerman. Still dead after all these years: interpreting the failure of general equilibrium theory. *Journal of Economic Methodology*, 9(2):119–139, June 2002.
- [5] K. Arrow, H. Block, and L. Hurwicz. On the stability of the competitive equilibrium, ii. *Econometrica*, 27(1):82–109, January 1959.
- [6] K. Arrow and F. Hahn. *General Competitive Analysis*. Holden-Day, 1971.
- [7] K. Arrow and L. Hurwicz. On the stability of the competitive equilibrium, i. *Econometrica*, 26(4):522–552, October 1958.
- [8] K. Arrow and M. Mcmanus. A note on dynamic stability. *Econometrica*, 26(3):448–454, July 1958.
- [9] T. Basar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, London and San Diego, second edition, 1995.
- [10] A. Berman and R. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, 1979.
- [11] D. Bertsekas and J. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10:627–642, 2000.

- [12] L. S. Brakmo and L. L. Peterson. TCP Vegas: End to end congestion avoidance on a global Internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–1480, 1995.
- [13] D. Brown and C. Shannon. Uniqueness, stability, and comparative statics in rationalizable walrasian markets. *Econometrica*, 68(6):1529–1539, November 2000.
- [14] D. Carlson. A class of positive stable matrices. *Journal of Research of the National Bureau of Standards–B. Mathematical Sciences*, 78(1):1–2, March 1974.
- [15] C. Casetti, M. Gerla, S. Mascolo, M. Sanadidi, and R. Wang. TCP Westwood: End-to-end congestion control for wired/wireless networks. *Wireless Networks Journal*, 8(5):467–479, September 2002.
- [16] G. Coxson. The p-matrix problem is co-np complete. *Mathematical Programming*, 64:173–178, 1994.
- [17] G. Debreu. *Theory of Value*. New York:Wiley, 1959.
- [18] G. Debreu. Economies with a finite set of equilibria. *Econometrica*, 38(3):387–392, May 1970.
- [19] G. Debreu. Economic theory in the mathematical mode. *Nobel Memorial Lecture*, December 1983.
- [20] G. Debreu and I. Herstein. Nonnegative square matrices. *Econometrica*, 21(4):597–607, October 1953.
- [21] E. Dierker. Two remarks on the number of equilibria of an economy. *Econometrica*, 40(5):751–753, September 1972.
- [22] E. Dierker and H. Dierker. The local uniqueness of equilibria. *Econometrica*, 40(5):867–881, September 1972.
- [23] N. Dukkupati, M. Kobayashi, R. Z. Shen, and N. McKeown. Processor sharing flows in the Internet. In *Proceedings of IWQoS*, June 2005.

- [24] A. Enthoven and K. Arrow. A theorem on expectations and the stability of equilibrium. *Econometrica*, 24(3):288–293, July 1956.
- [25] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.
- [26] F. Gantmacher. *The Theory of Matrices*. Chelsea Publishing Company, 1959.
- [27] L. Grieco and S. Mascolo. Performance evaluation and comparison of Westwood+, New Reno, and Vegas TCP congestion control. *ACM SIGCOMM Computer Communication Review*, 34(2):25–38, April 2004.
- [28] D. Hawkins and H. Simon. Note: Some conditions of macroeconomic stability. *Econometrica*, 17(3/4):245–248, July to October 1949.
- [29] D. Hershkowitz. Recent directions in matrix stability. *Linear Algebra and Its Applications*, 171:161–186, 1992.
- [30] J. Hicks. *Value and Capital*. London, 1939.
- [31] O. Holtz. Not all gkk τ -matrices are stable. *Linear Algebra and Its Applications*, 291:235–244, 1999.
- [32] V. Jacobson. Congestion avoidance and control. In *Proceedings of ACM Sigcomm*, August 1988.
- [33] R. Jain. A delay-based approach for congestion avoidance in interconnected heterogeneous computer networks. *Computer Communications Review*, 19(5):56–71, October 1989.
- [34] R. Jain, W. Hawe, and D. Chiu. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report DEC TR-301, Digital Equipment Corporation, September 1984.
- [35] C. Jin, D. X. Wei, and S. H. Low. FAST TCP: motivation, architecture, algorithms, performance. In *Proceedings of IEEE Infocom*, March 2004.

- [36] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
- [37] D. Katabi, M. Handley, and C. Rohrs. Congestion control for high bandwidth-delay product networks. In *Proceedings of ACM Sigcomm*, 2002.
- [38] A. Katok and B. Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1995.
- [39] F. P. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of Operations Research Society*, 49(3):237–252, March 1998.
- [40] H. Khalil. *Nonlinear Systems*. Prentice Hall, 2002.
- [41] S. Kunniyur and R. Srikant. End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Transactions on Networking*, 11(5):689 – 702, October 2003.
- [42] K. Kurata, G. Hasegawa, and M. Murata. Fairness comparisons between TCP Reno and TCP Vegas for future deployment of TCP Vegas. In *Proceedings of INET*, 2000.
- [43] A. Kuzmanovic. The power of explicit congestion notification. In *Proceedings of ACM Sigcomm*, August 2005.
- [44] S. H. Low. A duality model of tcp and queue management algorithms. *IEEE/ACM Transactions on Networking*, 11(4):525–536, August 2003.
- [45] S. H. Low, L. L. H. Andrew, and B. P. Wydrowski. Understanding XCP: equilibrium and fairness. In *Proceedings of IEEE Infocom*, March 2005.
- [46] S. H. Low and D. Lapsley. Optimization flow control, i: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, December 1999.
- [47] S. H. Low, L. Peterson, and L. Wang. Understanding Vegas: a duality model. *Journal of ACM*, 49(2):207–235, March 2002.

- [48] J. Marsden and M. Hoffman. *Elementary Classical Analysis*. W. H. Freeman, New York, second edition, 1993.
- [49] A. Mas-Colell. *The Theory of General Economic Equilibrium. A Differentiable Approach*. Cambridge University Press, 1985.
- [50] A. Mas-Colell. The future of general equilibrium. *Spanish Economic Review*, 1:207–214, 1999.
- [51] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [52] C. Meyer. *Matrix Analysis and Applied Linear Algebra*. SIAM, 2000.
- [53] J. Milnor. *Topology from the Differentiable Viewpoint*. The University Press of Virginia, Charlottesville, 1972.
- [54] J. Mo, R. La, V. Anantharam, and J. Walrand. Analysis and comparison of TCP Reno and Vegas. In *Proceedings of IEEE Infocom*, March 1999.
- [55] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, October 2000.
- [56] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [57] I. Pearce and J. Wise. On the uniqueness of competitive equilibrium: part i, unbounded demand. *Econometrica*, 41(5):817–828, September 1973.
- [58] I. Pearce and J. Wise. On the uniqueness of competitive equilibrium: part ii, bounded demand. *Econometrica*, 42(5):921–932, September 1974.
- [59] K. Ramakrishnan, S. Floyd, and D. Black. The addition of explicit congestion notification (ECN) to IP. RFC 3168, Internet Engineering Task Force, September 2001. <http://www.ietf.org/rfc/rfc3168.txt>.
- [60] A. Rantzer. A dual to Lyapunov’s stability theorem. *Systems and Control Letters*, 42(3):161–168, 2001.

- [61] L. Rizzo. Dummynet. http://info.iet.unipi.it/~luigi/ip_dummynet/.
- [62] T. Roughgarden and E. Tardos. How bad is selfish routing. *Journal of the ACM*, 49(2):236–259, 2002.
- [63] P. Samuelson. The relation between hicksian stability and true dynamic stability. *Econometrica*, 12(3/4):256–257, July to October 1944.
- [64] P. Samuelson. *Foundations of Economic Analysis*. Harvard University Press, 1947.
- [65] H. Scarf. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3):157–172, September 1960.
- [66] H. Schneider. An inequality for latent roots applied to determinants with dominant principal diagonal. *Journal of the London Mathematical Society*, 28:9–20, 1953.
- [67] A. Simsek, A. Ozdaglar, and D. Acemoglu. Generalized Poincare-Hopf theorem for compact nonsmooth regions. *Submitted for publication*, 2005.
- [68] S. Smale. Dynamics in general equilibrium theory. *The American Economic Review*, 66(2):288–294, 1976.
- [69] H. Smith. *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, 1995.
- [70] S. Sternberg. *Lectures on Differential Geometry*. Prentice Hall, 1964.
- [71] K. Tan, J. Song, Q. Zhang, and M. Sridharan. A compound TCP approach for high-speed and long distance networks. In *Proceedings of IEEE Infocom*, April 2006.
- [72] A. Tang, J. Wang, S. Hedge, and S. H. Low. Equilibrium and fairness of networks shared by TCP Reno and Vegas/FAST. *Telecommunication Systems*, 30(4):417–439, December 2005.
- [73] A. Tang, J. Wang, and S. H. Low. Is fair allocation always inefficient. In *Proceedings of IEEE Infocom*, March 2004.

- [74] A. Tang, J. Wang, and S. H. Low. Counter-intuitive throughput behaviors in networks under end-to-end control. *IEEE/ACM Transactions on Networking*, 14(2):355–368, April 2006.
- [75] A. Tang, J. Wang, S. H. Low, and M. Chiang. Equilibrium of heterogeneous congestion control protocols. *Caltech CS Technical Report: 2005.005*, April 2005.
- [76] A. Tang, J. Wang, S. H. Low, and M. Chiang. Network equilibrium of heterogeneous congestion control protocols. In *Proceedings of IEEE Infocom*, March 2005.
- [77] A. Tang, D. Wei, S. H. Low, and M. Chiang. Bandwidth allocation for networks with heterogeneous congestion control protocols. *Submitted for publication*, 2006.
- [78] H. Varian. A third remark on the number of equilibria of an economy. *Econometrica*, 43(5/6):985–986, September to November 1975.
- [79] Z. Wang and J. Crowcroft. Eliminating periodic packet losses in the 4.3-Tahoe BSD TCP congestion control algorithm. *ACM Computer Communications Review*, 22(2), April 1992.
- [80] D. Wei, C. Jin, S. H. Low, and S. Hegde. FAST TCP: motivation, architecture, algorithms, performance. *To appear in IEEE/ACM Transactions on Networking*, 2007.
- [81] B. Wydrowski, L. L. H. Andrew, and M. Zukerman. Maxnet: A congestion control architecture for scalable networks. *IEEE Communications Letters*, 7(10):511–513, October 2003.
- [82] B. Wydrowski and M. Zukerman. Maxnet: A congestion control architecture. *IEEE Communications Letters*, 6(11):512–514, November 2002.
- [83] Y. Xia, L. Subramanian, I. Stoica, and S. Kalyanaraman. One more bit is enough. In *Proceedings of ACM Sigcomm*, August 2005.
- [84] H. Yaiche, R. R. Mazumdar, and C. Rosenberg. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, 8(5):667–678, 2000.