DYNAMIC SELFISH ROUTING

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der Rheinisch-Westfälischen Technischen Hochschule Aachen zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Diplom-Informatiker

SIMON FISCHER

aus Hagen

Berichter: Universitätsprofessor Dr. Berthold Vöcking
Universitätsprofessor Dr. Petri Mähönen
Tag der mündlichen Prüfung: 13. Juni 2007

Dynamic Selfish Routing

SIMON FISCHER

Aachen • June 2007

Abstract

This thesis deals with dynamic, load-adaptive rerouting policies in game theoretic settings. In the Wardrop model, which forms the basis of our dynamic population model, each of an infinite number of agents injects an infinitesimal amount of flow into a network, which in turn induces latency on the edges. Each agent may choose from a set of paths and strives to minimise its sustained latency selfishly. Population states which are stable in that no agent can improve its latency by switching to another path are referred to as *Wardrop equilibria.*

A variety of results in this model have been obtained recently. Most of these revolve around the *price of anarchy*, which measures the degradation of performance due to the selfish behaviour of the agents as compared to a centrally optimised solution. Most of these analyses consider the notion of Wardrop equilibria as a solution concept of selfish routing games, but disregard the question of *how* such an equilibrium can be attained in the first place. In fact, common game theoretic assumptions needed for the motivation of equilibria in general strategic games are not satisfied for routing games in large networks, the Internet being the prime example. These assumptions comprise accurate knowledge of the network and its latency functions as well as unbounded rationality and reasoning capabilities. The question of how Wardrop equilibria can be attained by a population of selfish agents is the central topic of this thesis.

Our first approach is inspired by evolutionary game theory. We show that Wardrop equilibria actually satisfy a stronger stability criterion, called evolutionary stability, which can be motivated by milder assumptions. We model a population of agents following some very simple randomised selfish rules allowing them to exchange their path for a better one from time to time. An agent samples another path according to some probability distribution and possibly exchanges its current path for the new one with a probability that increases with the latency gain. The behaviour of such a population over time can be described in terms of a system of differential equations. For a concrete choice of probability distributions involved in the above rules, these differential equations take the form of the so-called

replicator dynamics. As a first result, we show that a population following this dynamics converges towards the set of Wardrop equilibria.

This convergence result implicitly assumes perfect information in that rerouting decisions made by other agents are observable by other agents immediately. In communication networks, however, such rerouting decisions may be observed only with a delay. In both theory and practise, it is known that rerouting decisions based on stale information may lead to undesirable oscillation effects which seriously harm performance. We consider an extension of our dynamic population model in which information is updated only at intervals of a given length. We show that, despite of this, convergence to Wardrop equilibria is still possible. For any class of latency functions with bounded slope and any finite update period length, policies from a large class converge towards Wardrop equilibria provided that they satisfy a certain smoothness condition. This condition requires the migration rate to be reduced by a factor that is reciprocal to the maximum slope and the update period length.

Finally, we show that Wardrop equilibria can be approached quickly. This is an issue of particular importance if the network or the flow demands themselves may change over time. To measure the speed of convergence, we consider the time to reach approximate Wardrop equilibria. We show that by applying a clever sampling technique it is possible to reach an approximate Wardrop equilibrium in time polynomial in the approximation quality and the representation length of the network, e. g., for positive polynomial latency functions in coefficient representation. In particular, our bounds depend on the maximum slope of the latency functions only logarithmically, which improves over earlier results in similar models which depend linearly on this parameter. We show that the crucial parameter that limits the speed of convergence is not the slope, but rather the elasticity of the latency functions. This parameter is closely related to the degree of polynomials.

Based on these positive theoretical results, we design a dynamic traffic engineering protocol which we evaluate by simulations. Our protocol splits traffic bound for the same destination among alternative next-hop routers and continuously adjusts the splitting ratios. The simulation framework involves significantly more details than the Wardrop model does. For example, our simulations feature a full TCP implementation at packet level and include realistic HTTP workload generated on the basis of realistic statistical properties of Web users. The simulation results are in good correspondence with theory. We see that our protocol in fact converges quickly and significantly improves the throughput of an autonomous system.

Acknowledgements

First of all, I thank Berthold Vöcking for being my advisor, not only with respect to this thesis, for his clear-sightedness, and for always knowing what is important and what is not. I thank Petri Mähönen for his interest in this work and for acting as a co-referee.

I thank Harald Räcke, Nils Kammenhuber, and Anja Feldmann for the productive collaboration and many inspiring discussions. I am particularly grateful to Nils for introducing me to SSFNet, though this gratefulness is kind of ambivalent.

I also thank the whole team of Computer Science 1 at Aachen for the nice working atmosphere, many interesting discussions, and also for assigning me to exercise six of the BuK 2007 test. I am in debt to Lars Olbrich for reading through early versions of some of the proofs.

Finally, I thank Gisela for her substantial support, in particular during the last four weeks of writing up this thesis.

This research was supported by the DFG Schwerpunktprogramm 1126 "Algorithmik großer und komplexer Netzwerke" and partially by the EU within the 6th Framework Programme under contract 001907 (DELIS).

Table of Contents

List of Figures

Notation

$\mathcal{C}$	the set of commodities, where commodity $C_i = (s_i, t_i, r_i)$
D	maximum length of a path
Δ_n	the n-dimensional simplex $\Delta_n = \left\{ \mathbf{x} \in \mathbb{R}^n_{\geq 0} \mid \sum_{i \in [n]} x_i = 1 \right\}$
$\mathbf{e}_k$	the unit vector in dimension k
$\mathcal{F}_\Gamma$	the set of feasible flow vectors; $\mathcal{F}_\Gamma^{NE}$ is the set of Wardrop equilibria
f	path flow vector $f = (f_P)_{P \in \mathcal{P}}$; $f_e = \sum_{P \ni e} f_P$ is the flow on edge e
Γ	an instance $\Gamma = (G, (\ell_e)_{e \in E}, \mathcal{C})$ of the Wardrop routing game
$\mathrm{int}(\Delta)$	the interior of the simplex, i. e. the set $\{\mathbf{x} \in \Delta \mid x_i > 0 \; \forall i \in [n]\}$
k	the number of commodities
ℓ_e	latency function of edge e
L	average latency
$[n]$	the set $\{1, \ldots, n\}$
$\mathbb{N}$	the set of natural numbers (excluding 0)
Φ	Beckmann-McGuire-Winsten potential $\Phi(f) = \sum_{e \in E} \int_0^{f_e} \ell_e(u)\, du$
$\mathcal{P}_i$	paths between source s_i and sink t_i, $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$
$\mathbb{R}$	the set of real numbers; $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_+$ denote the non-negative and positive real numbers, respectively
$\mathrm{Supp}(\mathbf{x})$	the support of $\mathbf{x}$, i. e., the set $\{i \in [n] \mid x_i > 0\}$
$\dot{\mathbf{x}}$	time derivative of $\mathbf{x}(t)$, $\dot{\mathbf{x}}(t) = \frac{d}{dt}\mathbf{x}(t)$
$\langle \mathbf{x} \vert \mathbf{y} \rangle_\epsilon$	The convex combination $(1-\epsilon)\,\mathbf{x} + \epsilon\,\mathbf{y}$

Throughout this thesis, we use the following notation. The real, non-negative real, and positive real numbers are denoted $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, and $\mathbb{R}_+$, respectively. The set $\{1, \ldots, n\}$ is written as $[n]$, and for any set S, 2^S denotes the power set $2^S = \{T \mid T \subseteq S\}$. In n-dimensional space, the unit vector in dimension k is denoted $\mathbf{e}_k$. The dimension n will be clear from the context. By $\mathbf{x}_{|I}$ we denote the vector $\mathbf{x}$ restricted to the index set I.

An affine function is the sum of a linear function and a constant. The natural logarithm is denoted $\ln(\cdot)$, and the logarithm with base 2 is denoted $\log(\cdot)$.

Mathematical Preliminaries

For two functions $f, g : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$ we say that $f(x) = \mathcal{O}(g(x))$ as $x \to \infty$ if and only if there exist $x_0, c > 0$ such that $f(x) \leq c \cdot g(x)$ for all $x \geq x_0$. We say that $f(x) = \mathcal{O}(g(x))$ as $x \to 0$ if and only if there exist x_0, $c > 0$ such that $f(x) \leq c \cdot g(x)$ for all $x \leq x_0$. We will omit the phrases "as $x \to \infty$" and "as $x \to 0$" since it will be clear from context. If $f(x_1, \ldots, x_n)$ is a multivariate function, we require the above to hold for $x_i \geq x_{i,0}$ ($x_i \leq x_{i,0}$) for all $i \in [n]$ and positive constants $x_{i,0}$, $i \in [n]$. Correspondingly, we say $f(x) = \Omega(g(x))$ if and only if the above holds where the condition $f(x) \leq c \cdot g(x)$ is replaced with $f(x) \geq c \cdot g(x)$. If f satisfies $f(x) = \mathcal{O}(g(x))$ and $f(x) = \Omega(g(x))$, then $f(x) = \Theta(g(x))$. Finally, $f(x) = o(g(x))$ as $x \to a$, $a \in \{0, \infty\}$ if and only if $\lim_{x \to a} f(x)/g(x) = 0$.

A function $f(\cdot)$ is Lipschitz-continuous at x_0 if there exist constants $\epsilon > 0$ and $\alpha > 0$ such that for all $x < \epsilon$, $|f(x) - f(x_0)| \leq \alpha \cdot |x - x_0|$.

We will frequently make use of *Jensen's inequality* in its following form:

Theorem 0.1. *For a convex function $f : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$, numbers $x_i \in \mathbb{R}_{\geq 0}$ and weights $w_i \in \mathbb{R}_{\geq 0}$ it holds that*

$$\sum_i w_i \cdot f(x_i) \geq f\left(\frac{\sum_i w_i\, x_i}{\sum_i w_i}\right) \cdot \sum_i w_i \ .$$

In particular, we will consider the case that $f(x) = x^2$ and the w_i are probabilities, i. e., $\sum_i w_i = 1$. Then,

$$\sum_i w_i \cdot x_i^2 \geq \left(\sum_i w_i\, x_i\right)^2 \ .$$

CHAPTER 1

Introduction

The rise of the Internet as today's predominant communication medium has opened up new perspectives to computer science. The aspect in which the Internet differs most profoundly from many classical models studied in theoretical computer science is its lack of central coordination: Internet traffic does not lie in the hands of a single authority but is rather controlled by a large number of users acting, at the very best, selfishly, and possibly even maliciously, randomly, or faulty. This fact brings networking into the focus of game theory. In an influential paper, Roughgarden and Tardos [70] analyse routing problems in Internet-like networks using a model originally introduced by Wardrop in 1952 [83] as a model of road traffic. In this model, fixed flow demands are to be routed between source-sink-pairs, thus inducing load-dependent latencies on the links.

A classical objective in this scenario is to find a flow that optimises the overall performance of the network, e. g., by minimising the average latency. In the game theoretic setting, however, we assume that the flow is composed of infinitesimal contributions of an infinite number of selfish agents each of which strives for minimising its own sustained latency. We ask for an allocation of routes that is stable in that no agent has an incentive for deviating from this allocation. Such an allocation is called a *Nash equilibrium*. Compared to a centrally optimised flow, the overall performance of such an equilibrium may be worse. This degradation of performance is measured by the *price of anarchy* which has been studied extensively, e. g., in [6, 19, 21, 66, 67, 68, 70, 71]. Several approaches to reduce the price of anarchy have been studied, e. g., by introducing taxes [20, 35] or by controlling a subset of the agents centrally [65].

All of the above studies focus on static properties of Nash equilibria and are based on the prerequisite that selfish behaviour in network routing games results in such an equilibrium. Whereas for all reasonable models of selfish behaviour a population of agents should be stable once it has reached a Nash equilibrium, it is not clear *how* a population of agents can attain an equilibrium in the first place. Game theory uses several strong assumptions to motivate Nash equilibria for general strategic games, including complete and accurate knowledge about the payoff function (in our case given by the latency dependence of the links) as well as unbounded rationality and reasoning capabilities of the agents. Arguably, none of these assumptions are realistic in real-world applications, the Internet being the prime example.

The aim of this thesis is to circumvent these assumptions and to shed some light on processes by which a population of selfish agents is capable of attaining a Nash equilibrium in a simple and distributed fashion relying on very mild assumptions only. We believe that analyses of static properties of Nash equilibria are meaningful only if the limit behaviour of some natural process of the above kind results in a Nash equilibrium. Furthermore, it is desirable that the Nash equilibrium is approached quickly.

To a large extent, our dynamic population model is inspired by evolutionary game theory. Before presenting an outline of our results in Section 1.4, we will formally introduce the models. After a brief introduction to game theoretic concepts and notation in Section 1.1 we present the dynamic extensions and refinements of evolutionary game theory in Section 1.2. The Wardrop model is formally introduced in Section 1.3.

1.1 Foundations of Game Theory

In the following, we assume a basic knowledge of (classical) game theory. To establish notation, we will briefly introduce the basic concepts. For a more elaborate introduction, we refer to standard text books [53, 56, 80].

A strategic *symmetric two-player game* can be described by an $n \times n$-*payoff-matrix* $\mathbf{A} = (a_{ij})_{i,j \in [n]}$. The set of *pure strategies* of both players is given by the set $[n] = \{1, \dots, n\}$. The term *symmetric* refers to the fact that both players have the same strategy space and payoff matrix $\mathbf{A}$ as opposed to the general case where we have an individual strategy space and payoff matrix for each player. Since only symmetric games qualify for our subsequent considerations, we restrict ourselves to the symmetric case here.

Suppose that the two players decide for strategy $i \in [n]$ and $j \in [n]$, respectively. The *payoff* earned by player 1 when playing pure strategy $i \in [n]$ against the opponent's pure strategy $j \in [n]$ is given by a_{ij} and the opponent's payoff is given by a_{ji}. A *mixed strategy* $\mathbf{x} = (x_i)_{i \in [n]}$ is a probability distribution over the strategy space $[n]$. The set of all mixed strategies is the simplex

$$\Delta_n = \left\{ \mathbf{x} \in \mathbb{R}^n_{\geq 0} \mid \sum_{i \in [n]} x_i = 1 \right\} \quad .$$

The expected payoff of player 1 playing the pure strategy i against a mixed strategy $\mathbf{x}$ is then $(\mathbf{A}\mathbf{x})_i = \sum_{j \in [n]} a_{ij} \cdot x_j$. Finally, the expected payoff of player 1 playing the mixed strategy $\mathbf{y}$ against the mixed strategy $\mathbf{x}$ is $\mathbf{y}^\top \mathbf{A}\mathbf{x}$. Let $\mathbf{e}_k$ denote the k-th unit vector in n-dimensional space. For any mixed strategy $\mathbf{x}$ which assigns positive probability only to the pure strategy k, i. e., $\mathbf{x} = \mathbf{e}_k$, we identify $\mathbf{x}$ with the pure strategy k.

The central question of game theory is: How should two strategic players involved in the game behave, given a payoff matrix $\mathbf{A}$? Clearly, a rational player knowing the payoff matrix should never decide for a strategy $\mathbf{x}$ if there exists another strategy $\mathbf{y_x}$ such that for every strategy of the opponent $\mathbf{z}$, $\mathbf{y_x}^\top \mathbf{A}\mathbf{z} > \mathbf{x}^\top \mathbf{A}\mathbf{z}$. We say that such a strategy $\mathbf{y_x}$ *strictly dominates* strategy $\mathbf{x}$. Hence, under this rationality assumption we may remove any pure dominated strategy from the game without affecting the outcome. Once this is done, other pure strategies may become dominated. Iterating this process, we can remove all *iteratively strictly dominated strategies* from the game. Note that the first iteration of this process merely requires the first player to know her own preferences. For subsequent iterations, every player must assume that the payoff matrix truly reflects the preferences of the other player and that the other player also acts rational in this sense. Whereas some games are *strictly dominance solvable* in that this process leaves just a single strategy to each player, no dominated strategies have to exist in general.

Now, let us assume that player 1 knows the strategy $\mathbf{y}$ chosen by player 2. Then, player 1 will choose a strategy $\mathbf{x}$ that maximises her payoff $\mathbf{x}^\top \mathbf{A}\mathbf{y}$. Any such strategy $\mathbf{x}$ that maximises the payoff $\mathbf{x}^\top \mathbf{A}\mathbf{y}$ against a fixed strategy $\mathbf{y}$ is called a *best response* to $\mathbf{y}$. The set of best responses to $\mathbf{y}$ is denoted by $\beta(\mathbf{y})$. If, in turn, player 2 knows the first player's strategy $\mathbf{x}$ in advance, player 2 will choose a best reply to $\mathbf{x}$. Consequently, a pair of strategies $(\mathbf{x}, \mathbf{y})$ will be stable only if $\mathbf{x}$ and $\mathbf{y}$ are mutual best responses. This is one way to define Nash equilibria, the predominant solution concept for strategic games.

Definition 1.1 (Nash equilibrium). *Given a symmetric two-player game* $\mathbf{A}$*, a (mixed) strategy pair* $(\mathbf{x}, \mathbf{y})$ *is a* Nash equilibrium *for* $\mathbf{A}$ *if* $\mathbf{x} \in \beta(\mathbf{y})$ *and* $\mathbf{y} \in \beta(\mathbf{x})$*. The (mixed) strategy* $\mathbf{x}$ *is a* symmetric Nash equilibrium *for* $\mathbf{A}$ *if* $\mathbf{x} \in \beta(\mathbf{x})$*. The set of symmetric Nash equilibria is denoted by* Δ^{NE}*.*

Thus, a Nash equilibrium is a pair of strategies, such that no player has an incentive to deviate from her strategy unilaterally. Note that symmetric games may possess Nash equilibria that are not symmetric themselves. For example, the symmetric game

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

possesses the pure Nash equilibria $(\mathbf{e}_1, \mathbf{e}_2)$ and $(\mathbf{e}_2, \mathbf{e}_1)$ which are not symmetric.

It is not obvious that Nash equilibria always exist, and, in fact, it is easy to check that *pure* equilibria do not always exist. Consider the children's game *Rock-Scissors-Paper* in which each of the two players simultaneously shows one out of three symbols (rock, scissors, and paper) where rock beats scissors, scissors beat paper, and paper beats rock. The winner earns one point from the loser, and two equal symbols are a draw, giving zero payoff to both. The payoff matrix of this game is

$$\mathbf{A}_{PRS} = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} .$$

It is easy to check that for any pure strategy, $i \in \{1, 2, 3\}$, the best reply is uniquely determined, i. e., $\beta(i)$ is a singleton. Furthermore, $i \notin \beta(\beta(i))$. Consequently, no pair of pure strategies can be a Nash equilibrium. A similar argument holds for any biased mixed strategy. Hence, the unbiased strategy $(1/3, 1/3, 1/3)$ is the only Nash equilibrium of this game. It can be shown that, if we allow the players to randomise their strategic choices, a Nash equilibrium always exists.

Theorem 1.1. *For any strategic game, a Nash equilibrium exists. For any symmetric two-player game, $\Delta^{NE} \neq \emptyset$.*

We have introduced Nash equilibria as strategy pairs which are stable in the sense that no player has an incentive to deviate from her strategy given a belief about the opponent's strategy. Thus, if we assume that strategies are proposed to the players by some external entity like society, tradition, or a coordinator, then Nash equilibria are those strategies which are compatible with the players' incentives. This motivation, however, does not explain how players can attain such a Nash equilibrium by themselves.

For another motivation assume that players have complete and accurate knowledge about the game, i. e., the payoff matrix. Furthermore, assume that

1. both players are rational,

2. both players know (1),

3. both players know (2),

   ...

i. both players know $(i-1)$, etc.

   ...

The latter chain of statements is referred to as *common knowledge of unbounded rationality*. If the players' capabilities of reasoning are unlimited, Nash equilibria are the only strategies that players can identify as coherent with the above assumptions. Still, this does not answer the question *which* Nash equilibrium players should agree upon in case of the coexistence of several. Consider the game

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

with its two symmetric Nash equilibria $\mathbf{e}_1$ and $\mathbf{e}_2$. Any argument (that is independent of the numbering of the strategies) that would favour one of the Nash equilibria over the other could also be used to favour the other due to the symmetry of the game.

We consider a third approach for the motivation of Nash equilibria. We view equilibria as the limit outcome of a process in which players can learn how to perform well in a game over time. It will turn out that under certain conditions, such processes lead to Nash equilibria (with particular properties) in the long run. We will elaborate on this approach in the next section.

1.2 Evolutionary Game Theory: A Guided Tour

In the preceding section we have used three major assumptions for the motivation of Nash equilibria. These assumptions include complete and accurate knowledge of the payoff matrix, common knowledge of unbounded rationality, and unlimited reasoning capabilities. In many settings in which strategic decisions have to be made, the impact of the own and other participants' decisions on the payoff can hardly be predicted. The objective, and hence the payoff of the other participants may not even be publicly known. Furthermore, participants in such a setting may not make the effort to make any such predictions at all, and if they do, they may even fail. This rules out most of the knowledge and rationality assumptions. Evolutionary approaches offer solution concepts of strategic games that rely on assumptions that are significantly milder than the ones presented above.

In the following two sections we will discuss two approaches to evolutionary game theory which are largely independent of each other. The first approach describes dynamic processes by which populations of individuals can come to an equilibrium. The second focuses on equilibrium refinements based on the resistance of populations to mutant population entrants. Finally, in Section 1.2.3 we will see how these approaches relate to each other. For a comprehensive introduction to evolutionary game theory, see, e. g., [84] or [42].

1.2.1 The First Approach: Evolutionary Game Dynamics

A large part of the inspiration for evolutionary game theory comes from biology. We consider a large population of individuals of the same type. Each individual exhibits a certain behaviour, corresponding to a pure strategy of the underlying game (to which it may be genetically or otherwise programmed). The individuals are repeatedly paired at random to play a symmetric game against each other. The expected payoff is then interpreted as the individual's *fitness* within the population.

We model the dynamic behaviour of such a population in a continuous fashion. We assume that the number of individuals is infinite and denote the population share programmed to play pure strategy i by x_i. Then, the vector $\mathbf{x} = (x_i)_{i \in [n]}$ is formally equivalent to a mixed strategy. Also, the expected payoff of playing against an agent randomly chosen from $\mathbf{x}$ is equal to the expected payoff playing against the mixed strategy $\mathbf{x}$.

We consider the population vector $\mathbf{x}$ as a function of time $\mathbf{x} : \mathbb{R}_{\geq 0} \mapsto \Delta$ where $\mathbf{x}(t)$ is the population state at time t. Although our motivation for the population dynamics to be derived below will be a discrete stochastic process, we can describe the population dynamics as a dynamical system in terms of differential equations. To that end, by the law of large numbers, we consider the limit behaviour of this process as the number of individuals approaches infinity, the so-called *fluid limit*.

We denote the time derivative of a population share $x_i(t)$ by $\dot{x}_i(t) = \frac{d}{dt} x_i(t)$ and suppress the argument t. For simplicity, we normalise the payoff matrix such that $\min_{i,j \in [n]} a_{ij} = 0$ and $\max_{i,j \in [n]} a_{ij} = 1$. Nash equilibria are invariant under this transformation.

1. *Fitness-proportional replication.* Suppose that individuals reproduce at a certain rate which is proportional to their fitness $(\mathbf{A}\mathbf{x})_i$. Then, the probability that an individual programmed to play strategy i is chosen for replication is

$$\frac{x_i \cdot (\mathbf{A}\mathbf{x})_i}{\sum_{j \in [n]} x_j \cdot (\mathbf{A}\mathbf{x})_j} = \frac{x_i \cdot (\mathbf{A}\mathbf{x})_i}{\mathbf{x}^\top \mathbf{A}\mathbf{x}} \quad .$$

If individuals breed true, i. e., the offspring of an individual is programmed to play the same strategy as the parent, and the reproduction and birth rates are 1, then the rate at which the population share using strategy $i \in [n]$ increases or decreases is

$$\dot{x}_i = \frac{x_i \cdot (\mathbf{A}\mathbf{x})_i}{\mathbf{x}^\top \mathbf{A}\mathbf{x}} - x_i = \frac{1}{\mathbf{x}^\top \mathbf{A}\mathbf{x}} \cdot x_i \cdot \left((\mathbf{A}\mathbf{x})_i - \mathbf{x}^\top \mathbf{A}\mathbf{x}\right) \ .$$

2. *Imitation of successful players.* In economic settings we do not consider individuals that reproduce, but rather individuals that change their strategy over time, continuously adapting their behaviour to the current population state. Suppose that players review their strategy at certain time steps and assume that these time steps are the arrival times of a Poisson process with arrival rate 1. Whenever activated, a player samples another player at random and observes this player's average payoff. It then imitates this player's strategy with a probability that is proportional to the payoff difference (and stays with her old strategy if this difference is negative). For simplicity, assume that the constant of proportionality is 1. Summing up inflow and outflow,

$$\begin{aligned} \dot{x}_i &= -\sum_{j:(\mathbf{A}\mathbf{x})_j \geq (\mathbf{A}\mathbf{x})_i} x_i\, x_j ((\mathbf{A}\mathbf{x})_j - (\mathbf{A}\mathbf{x})_i) \\ &\quad + \sum_{j:(\mathbf{A}\mathbf{x})_j < (\mathbf{A}\mathbf{x})_i} x_j\, x_i ((\mathbf{A}\mathbf{x})_i - (\mathbf{A}\mathbf{x})_j) \\ &= x_i \cdot \left((\mathbf{A}\mathbf{x})_i - \mathbf{x}^\top \mathbf{A}\mathbf{x}\right) \ . \end{aligned}$$

Alternatively, we may assume that an individual imitates her opponent with probability 1 whenever the opponent's fitness is higher than her own, but the expected payoffs are observed with uncertainty or noise. The observed expected payoff for strategy j is then $(\mathbf{A}\mathbf{x})_j + X$ where X is some random variable. If X is uniformly distributed over $[0, 1]$, we obtain the above expression for $\dot{x}_i$ again.

3. *Imitation driven by dissatisfaction.* Note that in order to execute the previous imitation process, individuals need to observe other individuals' payoffs. In cases where this is not possible, one may consider the following process. Again, assume that agents are activated at Poisson rates. Whenever activated, an individual picks an aspiration level from the range $[0, 1]$ uniformly at random. Whenever a player falls short of this aspiration level, it decides to try another strategy by imitating another individual chosen uniformly at random. Then,

$$\begin{aligned} \dot{x}_i &= -x_i \cdot (1 - (\mathbf{A}\mathbf{x})_i) + \sum_{j \in [n]} x_j \cdot (1 - (\mathbf{A}\mathbf{x})_j) \cdot x_i \\ &= x_i \cdot \left((\mathbf{A}\mathbf{x})_i - \mathbf{x}^\top \mathbf{A}\mathbf{x}\right) \ . \end{aligned}$$

The above processes rely on different assumptions: For the first process, we assume that payoff correlates to a reproduction or survival rate due to an evolutionary selection mechanism that is not made explicit. The individuals themselves remain passive in this setting. For the second two, the individuals actively engage in the process. We assume that the payoff truly represents the individuals preferences, and we assume that players can observe this payoff. For the second process, individuals need to be able to observe the expected payoff of all strategies whereas for the third they need to observe only their own. The reasoning involved in the imitation process is limited to simple comparisons and a randomised decision and sampling step. In particular, it does not involve common knowledge of unbounded rationality, or of any rationality at all.

For various similar dynamics see, e. g. [2, 43, 84]. Remarkably, all of the three expressions for $\dot{x}_i$ derived above differ merely in a scalar factor $(\mathbf{x}^\top \mathbf{A}\mathbf{x})$ that is independent of the strategy i. The thus obtained system of differential equations is referred to as the *replicator dynamics*.

Definition 1.2 (replicator dynamics)**.** *For any given payoff matrix* $\mathbf{A}$ *and a positive Lipschitz-continuous scalar function* $\lambda : \Delta \mapsto \mathbb{R}_+$*, the* replicator dynamics *is given by the system of differential equations*

$$\begin{aligned} \dot{x}_i &= \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{A}\mathbf{x})_i - \mathbf{x}^\top \mathbf{A}\mathbf{x}) \\ \mathbf{x}(0) &= \mathbf{x_0} \ , \end{aligned} \tag{1.1}$$

for all $i \in [n]$*.*

The direction of the underlying vector field is independent of the scalar $\lambda(\mathbf{x})$, and so are the solution orbits. The function λ has an influence only on the speed at which this orbit is traversed. The most common choice for $\lambda(\mathbf{x})$ is the constant function $\lambda = 1$. Let us remark that even for this simple case, the system contains cubic terms and no analytical solution is known.

Let us mention some basic properties of the replicator dynamics:

1. *Existence of a unique solution.* The right hand side of Equation (1.1) is Lipschitz-continuous if $\lambda(\cdot)$ is Lipschitz-continuous. Hence, by the Picard-Lindelöf-Theorem (see, e. g. [16]) it possesses a unique solution for any boundary condition $\mathbf{x_0}$. For any $\mathbf{x_0} \in \Delta$, this solution is denoted by the function $\xi(\mathbf{x_0}, t)$. We suppress the argument $\mathbf{x_0}$ when it is clear from context or irrelevant.

2. *Invariance of the simplex.* It is intuitive that solutions $\xi(t)$ of Equation (1.1) stay inside the simplex for all $t \geq 0$. This can be verified by

showing that $\sum_{i\in[n]} \dot{x}_i = 0$:

$$\begin{aligned}\sum_{i\in[n]} \dot{x}_i &= \lambda(\mathbf{x}) \cdot \sum_{i\in[n]} x_i \cdot ((\mathbf{Ax})_i - \mathbf{x}^\top \mathbf{Ax}) \\ &= \lambda(\mathbf{x}) \cdot (\mathbf{x}^\top \mathbf{Ax} - \mathbf{x}^\top \mathbf{Ax}) \\ &= 0 \ .\end{aligned}$$

Furthermore, $\xi_i(t) \geq 0$ for all $t \geq 0$. If $\xi_i(0) \geq 0$ and $\xi_i(t) < 0$ for some $t > 0$, then $\xi_i(u) = 0$ for some $u \in [0, t)$ since ξ_i is continuous in t. Let $t_0 = \inf\{u \mid \xi_i(u) = 0\}$. Then, $\xi_i(t_0) = 0$ since the set over which the infimum is taken is closed and $\xi_i(u)$ is continuous in u which in turn implies that $\xi_i(u) = 0$ for all $u \geq t_0$.

3. *Invariance of the interior and the boundary.* We have just seen that $\xi_i(t_0) = 0$ for some t_0 implies $\xi_i(t) = 0$ for all $t > t_0$. Similarly, $\xi_i(t_0) > 0$ for some t_0 implies $\xi_i(t) > 0$ for all $t > t_0$. This implies that starting in the interior of the simplex, the solution orbit will stay inside the interior for all time. The same holds for the boundary. If a pure strategy is initially unused, the replicator dynamics will not detect it. We say that the replicator dynamics is not *innovative*.

4. *Dominated strategies vanish.* Consider a dominated strategy $i \in [n]$. Then there exists a strategy $\mathbf{y}$ dominating i, and by continuity of the payoff function and compactness of Δ there exists a constant $\delta > 0$ such that for every strategy $\mathbf{x} \in \Delta$, $\mathbf{y}^\top \mathbf{Ax} - (\mathbf{Ax})_i \geq \delta$. To see that $x_i(t) \to 0$, define $v_i(\mathbf{x}) = \ln(x_i) - \sum_{j=1}^n y_j \ln(x_j)$ for $\mathbf{x} \in \mathrm{int}(\Delta)$. Observe that every term in the sum is negative and hence the contribution of the sum is positive. Now,

$$\begin{aligned}\dot{v}_i(\mathbf{x}) &= \frac{\dot{x}_i}{x_i} - \sum_{j=1}^n \frac{y_j \, \dot{x}_j}{x_j} \\ &= ((\mathbf{Ax})_i - \mathbf{x}^\top \mathbf{Ax}) - \sum_{j=1}^n y_j \cdot ((\mathbf{Ax})_j - \mathbf{x}^\top \mathbf{Ax}) \\ &= (\mathbf{Ax})_i - \mathbf{y}^\top \mathbf{Ax} \\ &\leq -\delta \ .\end{aligned}$$

Consequently for $\mathbf{x_0} \in \mathrm{int}(\Delta)$, $v_i(\mathbf{x}) \to -\infty$ implying $x_i(t) \to 0$ as $t \to \infty$. Once x_i has dropped beneath a threshold, other strategies may become dominated. Thus, iteratively dominated strategies are eliminated subsequently, i. e., $x_i \to 0$ as $t \to \infty$ if i is iteratively strictly dominated and $\mathbf{x_0} \in \mathrm{int}(\Delta)$. For a formal argument, see [72].

5. *Payoff monotonicity.* For any pure strategy i, it is easy to see that higher payoff implies a larger growth rate. More formally, let $g_i(\mathbf{x}) = \dot{x}_i/x_i$ denote the growth rate of the population share programmed to play the pure strategy i. Then, $g_i(\mathbf{x}) > g_j(\mathbf{x})$ if and only if $(\mathbf{Ax})_i > (\mathbf{Ax})_j$, we say that $\mathbf{g}(\cdot)$ is *payoff monotonic*. We can generalise this definition from pure strategies to (sub-)populations composed of individuals playing different pure strategies. We say that $\mathbf{g}(\cdot)$ is *aggregate payoff monotonic* if for any $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n_{\geq 0}$, $\mathbf{y} \cdot \mathbf{g}(\mathbf{x}) > \mathbf{z} \cdot \mathbf{g}(\mathbf{x})$ if and only if $\mathbf{y}^\top \mathbf{Ax} > \mathbf{z}^\top \mathbf{Ax}$. Aggregate payoff monotonicity requires successful subpopulations to grow faster than less successful ones. In [72] it is shown that any dynamics satisfying this property has the form of the replicator dynamics for some function $\lambda(\mathbf{x})$.

How can such a dynamical system help us in the study of equilibria of games? A natural stability concept of a dynamical system is the notion of a fixed point. A point $\mathbf{x}$ is a *fixed point* with respect to the replicator dynamics if $\dot{x}_i = 0$ for all $i \in [n]$, i. e., the right hand side of Equation (1.1) vanishes. It is easy to check that Nash equilibria are fixed points of the replicator dynamics: At a Nash equilibrium $\mathbf{x}$, for every strategy i either $(\mathbf{Ax})_i = \mathbf{x}^\top \mathbf{Ax}$ or $\mathbf{x}_i = 0$, both implying $\dot{x}_i = 0$. It is also easy to check that there are fixed points of the replicator dynamics that are not Nash equilibria: Nash equilibria of subgames obtained by removing some of the strategies are fixed points of the replicator dynamics, too. These are located on the boundary of the simplex.

Example 1. *Reconsider the Rock-Scissors-Paper Game. This game possesses a unique Nash equilibrium $\mathbf{x} = (1/3, 1/3, 1/3)$ in the interior of the simplex. Optimistically, one could hope that a population of individuals playing this game repeatedly should come to this equilibrium in the long run. If we apply the replicator dynamics to this game, this hope is disappointed. Figure 1.1(a) shows that for any initial condition other than the equilibrium itself, the solution orbit of the replicator dynamics is a closed loop around $\mathbf{x}$.*

Even worse, we may look at the perturbed game

$$\mathbf{A}_{PRS}(\delta) = \begin{pmatrix} \delta & 1 & -1 \\ -1 & \delta & 1 \\ 1 & -1 & \delta \end{pmatrix}$$

where a draw is rewarded a payoff of δ to both players. In this game, if $\delta > 0$, not only does the orbit fail to converge, it even approaches the boundary moving farther and farther away from $\mathbf{x}$ (Figure 1.1(b)). Only for $\delta < 0$ does the orbit converge towards the Nash equilibrium $\mathbf{x}$ (Figure 1.1(c)).

We see that a population state does not qualify as a state which can evolve in the long run just by being a Nash equilibrium alone. We therefore

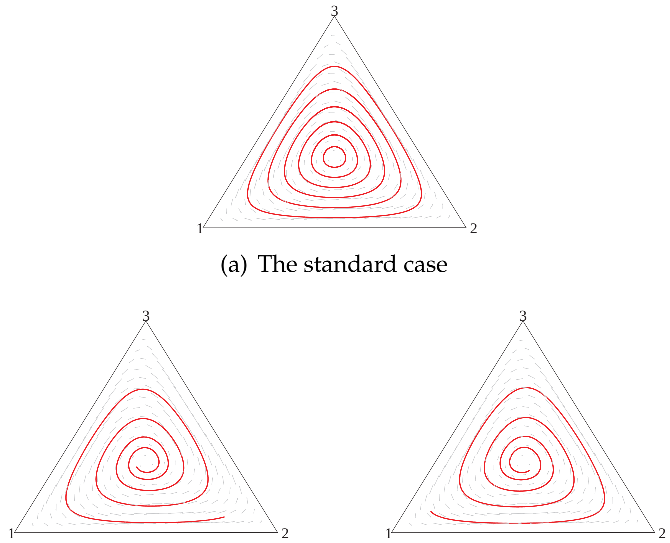

(a) The standard case

(b) Draw rewarded, $\delta > 0$

(c) Draw punished, $\delta < 0$

Figure 1.1: Solution orbits of the generalised Paper-Rock-Scissors game for different values of δ. Orbits are traversed counter-clockwise.

consider fixed points with stronger convergence properties. We introduce asymptotic stability not only for points, but for sets of points, since this will be useful later. For now, we may as well assume that the set X in the following definition is a singleton $\{\mathbf{x}\}$.

Definition 1.3 ((asymptotic) stability). *A closed set X is* stable *if every neighbourhood U of X contains a neighbourhood U' such that $\xi(\mathbf{x_0}, t) \in U$ for all $\mathbf{x_0} \in U'$ and $t \geq 0$. It is furthermore* asymptotically stable *if it is stable and there exists a neighbourhood U'' of X such that $\xi(\mathbf{x_0}, t)_{t\to\infty} \to X$ for all $\mathbf{x_0} \in U''$. Here, $\xi(\mathbf{x_0}, t) \to X$ means that* $\min_{\mathbf{x}\in X} ||\xi(\mathbf{x_0}, t), \mathbf{x}|| \to 0$.

The set of asymptotically stable Nash equilibria can be interpreted as a set of strategies that can be attained by a population of individuals using simple selfish steps that require only mild rationality assumptions. We have seen that this set is a proper subset of Δ^{NE}. However, it is not clear how a Nash equilibrium can be checked for asymptotic stability easily. To uphold the tension, we defer the answer to this question to Section 1.2.3

1.2.1.1 Other dynamics

A variety of dynamics is discussed in the literature, among which the replicator dynamics is the one most widely studied. For an overview, see [43]. For the replicator dynamics we have required only very limited reasoning capabilities of the agents. If we increase their reasoning capabilities slightly, assuming that they are able to compute or otherwise determine a best response with respect to the current population, we obtain the *best response dynamics*. Since the best response may not be unique and there is no obvious way of breaking ties this dynamics is a differential inclusion rather than a differential equation:

$$\dot{\mathbf{x}} \in \beta(\mathbf{x}) - \mathbf{x} \ . \tag{1.2}$$

As long as $\beta(\mathbf{x})$ is unique, it is a pure strategy, say i. Then the differential inclusion turns into a differential equation and is solved by

$$\mathbf{x}(t) = (1 - \exp(-t))\, \mathbf{e}_i + \exp(-t)\, \mathbf{x}(0) \ .$$

Now, for some $t_0 > 0$, $\beta(\mathbf{x}(t_0))$ may not be unique any more. Still, among all best replies to $\mathbf{x}(t_0)$ there exists one which is also a best response to itself. Denote such a strategy by $\mathbf{b}$. Since $\mathbf{b}$ is a best response to both $\mathbf{x}(t_0)$ and $\mathbf{b}$, it is also a best response to the convex combination $(1-\epsilon)\, \mathbf{x}(t_0) + \epsilon\, \mathbf{b}$ if $\epsilon \geq 0$ is small. Thus, we can construct a piecewise linear solution of Equation (1.2) by iterating this process.

An approximation of the best response dynamics is given by the *logit dynamics*, which has a unique solution:

$$\dot{x}_i = \frac{\exp((\mathbf{A}\mathbf{x})_i/\epsilon)}{\sum_{j \in [m]} \exp((\mathbf{A}\mathbf{x})_j/\epsilon)} - x_i$$

for a parameter $\epsilon > 0$. If ϵ is small, then the first term vanishes for all pure strategies that are no best replies.

The main advantage of the best response dynamics over the replicator dynamics is the fact that it is innovative in that it detects strategies that may be initially unused. We will discuss a drawback of the best response dynamics for our application in Chapter 3.

Another approach that makes the replicator dynamics innovative is to add a limited amount of mutation which enables the agents to explore the strategy space. Such dynamics are, e. g., considered in [38]. We will consider a similar approach in Chapter 4.

1.2.2 The Second Approach: Refined Solution Concepts

In this section we will formulate a refinement of Nash equilibria, so called *evolutionarily stable strategies*. Although the motivation for this refinement is based on a setting similar to the one described in the preceding section, it does not take into account any model of *how* the population evolves over time. Again, we consider a population of agents, each of which is programmed to play a certain strategy, which may now also be mixed. The population does not necessarily consist of an infinite number of agents.

Consider an incumbent strategy $\mathbf{x}$ played by all individuals of a population. In addition to the requirement of a Nash equilibrium which asks that an individual programmed to play $\mathbf{x}$ does not have an incentive to deviate from this strategy, we furthermore require that subpopulations playing any other strategy $\mathbf{y}$ have an incentive to return to $\mathbf{x}$, given that this population share is small enough. Such a strategy is said to be evolutionarily stable.

In economic settings we may consider such an equilibrium a convention. As an informal example, consider a state using the imperial system and a small subpopulation using the metric system. As long as communicating only within the subpopulation, the metric system is advantageous. However, within the entire population, the advantages of the metric system will outweigh the additional effort of unit conversion only if the subpopulation is large and hence this effort becomes infrequent. Hence, both the (pure) metric and the (pure) imperial system constitute evolutionarily stable strategies, where the metric system is robust against a larger portion of dissenters than the imperial is.

Formally, let $\mathbf{x}$ denote an incumbent strategy and let $\mathbf{y}$ denote a mutant strategy. If an ϵ-fraction of mutants invades the population, we obtain the post-entry population $\mathbf{x}_\epsilon = (1-\epsilon)\,\mathbf{x} + \epsilon\,\mathbf{y}$. We will write this as $\langle \mathbf{x}|\mathbf{y}\rangle_\epsilon = (1-\epsilon)\,\mathbf{x} + \epsilon\,\mathbf{y}$. In order to be evolutionarily stable, we require $\mathbf{x}$ to earn a higher average payoff in the post-entry population $\mathbf{x}_\epsilon$ than the mutant $\mathbf{y}$ does.

Definition 1.4 (evolutionarily stable strategy (ESS))**.** *A strategy* $\mathbf{x} \in \Delta$ *is said to be an* evolutionarily stable strategy *if for any strategy* $\mathbf{y} \in \Delta$ *there exists an* $\epsilon_\mathbf{y} \in (0,1)$ *such that for every* $\epsilon \in (0, \epsilon_\mathbf{y})$*,*

$$\mathbf{x}^\top \mathbf{A} \mathbf{x}_\epsilon > \mathbf{y}^\top \mathbf{A} \mathbf{x}_\epsilon \tag{1.3}$$

with $\mathbf{x}_\epsilon = \langle \mathbf{x}|\mathbf{y}\rangle_\epsilon$*. For any* $\mathbf{y}$*, the largest* $\epsilon_\mathbf{y}$ *satisfying the above condition is called its* invasion barrier $b(\mathbf{y})$*. The set of evolutionarily stable strategies is denoted by* Δ^{ESS}*.*

One may ask whether it is necessary that the parameter $\epsilon_\mathbf{y}$ in Definition 1.4 depends on the mutant $\mathbf{y}$. If these mutants could be arbitrarily far away from the equilibrium $\mathbf{x}$, this would obviously be the case. Since, however, the number of pure strategies is finite and the distance between $\mathbf{y}$ and $\mathbf{x}$ is bounded by the fact that both $\mathbf{x}$ and $\mathbf{y}$ are contained within Δ, we can choose $\epsilon_\mathbf{y}$ independent of $\mathbf{y}$. The strategies minimising $b(y)$ for any given direction from $\mathbf{x}$ are those on the boundary of the simplex.

Definition 1.5 (uniform invasion barrier)**.** *A strategy* $\mathbf{x} \in \Delta$ *has a* uniform invasion barrier $\bar{\epsilon} \in (0, 1)$ *if for all* $\mathbf{y} \in \Delta$, $b(\mathbf{y}) \geq \bar{\epsilon}$

Theorem 1.2. *A strategy* $\mathbf{x} \in \Delta^{ESS}$ *if and only if* $\mathbf{x}$ *has a uniform invasion barrier.*

For a formal proof, see [84]. The uniform invasion barrier $\bar{\epsilon}$ constitutes a lower bound on the size of the population that is necessary to make the population resistant against a single mutant individual. Consider a population with N individuals and assume that $N - 1$ individuals play the strategy $\mathbf{x} \in \Delta^{ESS}$ and a single individual mutates to play $\mathbf{y} \neq \mathbf{x}$. If $\bar{\epsilon} \geq 1/N$, then $\mathbf{y}$ performs worse in this population than the ESS strategy $\mathbf{x}$ does.

Still, this characterisation does not reveal how evolutionarily stable strategies relate to the well-known concept of Nash equilibria. It is easily verified that evolutionarily stable strategies are in fact Nash equilibria, but not all Nash equilibria are evolutionarily stable. In addition, evolutionary stability requires that the incumbent population $\mathbf{x}$ performs better against a mutant than this mutant performs against itself. This is guaranteed by the following nice characterisation of evolutionarily stable strategies.

Theorem 1.3. *A strategy* $\mathbf{x} \in \Delta$ *is evolutionarily stable if and only if* $\mathbf{x} \in \Delta^{NE}$ *and for all* $\mathbf{y} \in \beta(\mathbf{x})$, $\mathbf{y} \neq \mathbf{x}$, *it holds that* $\mathbf{x}^\top \mathbf{A}\mathbf{y} > \mathbf{y}^\top \mathbf{A}\mathbf{y}$.

Proof. First assume that $\mathbf{x}$ is evolutionarily stable. Assume, for the sake of contradiction, that $\mathbf{x}$ is not a Nash equilibrium, i. e., there exists another strategy $\mathbf{y}$ which earns higher payoff against $\mathbf{x}$ than $\mathbf{x}$ does, $\mathbf{y}^\top \mathbf{A}\mathbf{x} - \mathbf{x}^\top \mathbf{A}\mathbf{x} > 0$. Since for $\mathbf{x}_\epsilon = \langle \mathbf{x}|\mathbf{y}\rangle_\epsilon$ the difference $\mathbf{y}^\top \mathbf{A}\mathbf{x}_\epsilon - \mathbf{x}^\top \mathbf{A}\mathbf{x}_\epsilon$ is linear in ϵ, this term is positive for ϵ small enough, contradicting evolutionary stability of $\mathbf{x}$.

Second, assume that $\mathbf{x}$ is a Nash equilibrium, but there exists some $\mathbf{y} \in \beta(\mathbf{x})$, $\mathbf{y} \neq \mathbf{x}$ with $\mathbf{y}^\top \mathbf{A}\mathbf{y} \geq \mathbf{x}^\top \mathbf{A}\mathbf{y}$. Since $\mathbf{y}$ is a best reply to $\mathbf{x}$, also $\mathbf{y}^\top \mathbf{A}\mathbf{x} \geq \mathbf{x}^\top \mathbf{A}\mathbf{x}$. Consequently, for any convex combination $\mathbf{x}_\epsilon = \langle \mathbf{x}|\mathbf{y}\rangle_\epsilon$ it holds that $\mathbf{y}^\top \mathbf{A}\mathbf{x}_\epsilon \geq \mathbf{x}^\top \mathbf{A}\mathbf{x}_\epsilon$ contradicting evolutionary stability of $\mathbf{x}$.

Conversely, consider a Nash equilibrium $\mathbf{x}$ such that for any $\mathbf{y} \in \beta(\mathbf{x})$, $\mathbf{y} \neq \mathbf{x}$ we have $\mathbf{x}^\top \mathbf{A}\mathbf{y} > \mathbf{y}^\top \mathbf{A}\mathbf{y}$. Then, again, for any such $\mathbf{y} \in \beta(\mathbf{x})$, $\epsilon \in (0, 1)$, and $\mathbf{x}_\epsilon = \langle \mathbf{x}|\mathbf{y}\rangle_\epsilon$ we have $\mathbf{x}^\top \mathbf{A}\mathbf{x}_\epsilon > \mathbf{y}^\top \mathbf{A}\mathbf{x}_\epsilon$, since this holds for $\epsilon = 1$ and with weak inequality for $\epsilon = 0$. For any $\mathbf{y} \notin \beta(\mathbf{x})$, we have $\mathbf{y}^\top \mathbf{A}\mathbf{x} < \mathbf{x}^\top \mathbf{A}\mathbf{x}$ and Inequality (1.3) holds for $\epsilon > 0$ small enough by continuity. □

Let us apply this characterisation to an example.

Example 2. *Recall that the unique Nash equilibrium of the Rock-Scissors-Paper game is* $\mathbf{x} = (1/3, 1/3, 1/3)$. *This strategy earns zero payoff against any other pure or mixed strategy, and any strategy is a best reply. Similarly, any other strategy* $\mathbf{y}$ *earns zero payoff against itself, too, since the probability of winning equals the probability of loosing by symmetry. This implies that the Nash equilibrium is not evolutionarily stable.*

Consider the generalised Rock-Scissors-Paper game where a draw is rewarded an additional payoff of δ, $\mathbf{A}_{PRS}(\delta)$ from Example 1. For any δ, $\mathbf{x}$ is still the unique Nash equilibrium and any strategy is a best reply: For any $\mathbf{y}$, $\mathbf{x}^\top\mathbf{A}\mathbf{y} = \mathbf{y}^\top\mathbf{A}\mathbf{x} = \delta/3$ since the probability of a draw is $1/3$. The expected payoff of $\mathbf{y}$ against itself is

$$\begin{aligned}\mathbf{y}^\top\mathbf{A}\mathbf{y} &= \delta\cdot(y_1^2+y_2^2+y_3^2)+y_1y_2+y_2y_3+y_3y_1-y_1y_3-y_2y_1-y_3y_2\\ &= \delta\cdot(1-2\cdot(y_1+y_2-y_1^2-y_2^2-y_1y_2))\end{aligned}$$

where we have used that $y_3 = 1-y_1-y_2$. Observe that $y_1+y_2-y_1^2-y_2^2-y_1y_2 \leq 1/3$ with equality only for $y_1 = y_2 = y_3 = 1/3$. Therefore, if $\delta > 0$, $\mathbf{x}^\top\mathbf{A}\mathbf{y} < \mathbf{y}^\top\mathbf{A}\mathbf{y}$ for any $\mathbf{y} \neq \mathbf{x}$, and if $\delta < 0$, $\mathbf{x}^\top\mathbf{A}\mathbf{y} > \mathbf{y}^\top\mathbf{A}\mathbf{y}$ for any $\mathbf{y} \neq \mathbf{x}$. Consequently, the Nash equilibrium $\mathbf{x}$ is evolutionarily stable only for $\delta < 0$, that is, if a draw is punished.

The characterisations of evolutionary stability in terms of invasion barriers are slightly cumbersome in that they take into consideration three strategies: the equilibrium, a mutant, and a continuum of convex combinations of the above. Theorem 1.3 improves upon this, but it still requires to check for two conditions: the first order condition $\mathbf{x}^\top\mathbf{A}\mathbf{x} \geq \mathbf{y}^\top\mathbf{A}\mathbf{x}$, plus the second order condition $\mathbf{x}^\top\mathbf{A}\mathbf{y} > \mathbf{y}^\top\mathbf{A}\mathbf{y}$ in case that the first order condition only holds with equality. The last characterisation we present here circumvents this by simply requiring $\mathbf{x}$ to perform better against all nearby strategies, we say that $\mathbf{x}$ *locally superior*.

Definition 1.6. *A strategy $\mathbf{x}$ is* locally superior *if there exists a neighbourhood U of $\mathbf{x}$ such that $\mathbf{x}^\top\mathbf{A}\mathbf{y} > \mathbf{y}^\top\mathbf{A}\mathbf{y}$ for all $\mathbf{y} \neq \mathbf{x}$ in $U \cap \Delta$.*

Theorem 1.4. *A strategy $\mathbf{x}$ is locally superior if and only if $\mathbf{x} \in \Delta^{ESS}$.*

Proof. First, suppose that U is a neighbourhood of $\mathbf{x}$ such that $\mathbf{x}^\top\mathbf{A}\mathbf{y} > \mathbf{y}^\top\mathbf{A}\mathbf{y}$ for all $\mathbf{y} \in U \cap \Delta$, $y \neq \mathbf{x}$. Fix an arbitrary mutant $\mathbf{y} \in \Delta$ and observe that there exists some $\epsilon_{\mathbf{y}} \in (0,1)$ such that $\mathbf{x}_\epsilon = \langle \mathbf{x}|\mathbf{y}\rangle_\epsilon \in U$ for any $\epsilon \in (0, \epsilon_{\mathbf{y}})$. Now,

$$\mathbf{x}_\epsilon{}^\top\mathbf{A}\mathbf{x}_\epsilon = \epsilon\,\mathbf{y}^\top\mathbf{A}\mathbf{x}_\epsilon + (1-\epsilon)\mathbf{x}^\top\mathbf{A}\mathbf{x}_\epsilon$$

implying

$$\mathbf{x}_\epsilon{}^\top\mathbf{A}\mathbf{x}_\epsilon - \mathbf{x}^\top\mathbf{A}\mathbf{x}_\epsilon = \epsilon\cdot\left(\mathbf{y}^\top\mathbf{A}\mathbf{x}_\epsilon - \mathbf{x}^\top\mathbf{A}\mathbf{x}_\epsilon\right)\ .$$

Since $\mathbf{x}_\epsilon \in U$, the left hand side is negative, implying $\mathbf{x}^\top\mathbf{A}\mathbf{x}_\epsilon > \mathbf{y}^\top\mathbf{A}\mathbf{x}_\epsilon$ and $\mathbf{x} \in \Delta^{ESS}$.

Now assume that $\mathbf{x} \in \Delta^{ESS}$ and let $\bar{\epsilon}$ denote its uniform invasion barrier. Define

$$Z_{\mathbf{x}} = \{\mathbf{z} \in \Delta \mid z_i = 0 \text{ for some } i \in \mathrm{Supp}(\mathbf{x})\}$$

as the union of all boundary faces that do not contain $\mathbf{x}$. Furthermore, define

$$V = \{\mathbf{y} \in \Delta \mid \mathbf{y} = \langle \mathbf{x}|\mathbf{z_y}\rangle_\epsilon \text{ for some } \mathbf{z_y} \in Z_\mathbf{x} \text{ and } \epsilon \in [0, \bar{\epsilon})\} \ .$$

Observe that for any given $\mathbf{y}$, $\mathbf{z_y}$ in this definition is uniquely defined. Observe that V is closed and $\mathbf{x} \in V$. Hence, there exists a neighbourhood U of $\mathbf{x}$ such that $U \cap \Delta \subseteq V$. Fix any point $\mathbf{y} \in U \cap \Delta$, $\mathbf{y} \neq \mathbf{x}$. Since $\mathbf{y} \in V$ and $\mathbf{y}$ lies inside the invasion barrier of $\mathbf{z_y}$, $\mathbf{z_y}^\top \mathbf{Ay} < \mathbf{x}^\top \mathbf{Ay}$. Since trivially $\mathbf{x}^\top \mathbf{Ay} \leq \mathbf{x}^\top \mathbf{Ay}$ and $\mathbf{y}$ is a convex combination of $\mathbf{x}$ and $\mathbf{z_y}$, this implies that $\mathbf{y}^\top \mathbf{Ay} < \mathbf{x}^\top \mathbf{Ay}$. □

How many evolutionarily stable states can a game possess and how do they relate to each other? We will see that, if we know one evolutionarily stable state $\mathbf{x}$, we can be sure that certain parts of the simplex, those that conflict with $\mathbf{x}$ in a precise sense, do not contain further evolutionarily stable states. We can illustrate this with another example.

Example 3. *Consider the payoff matrix*

$$\begin{pmatrix} 1 & 2 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix} \ .$$

This game possesses three Nash equilibria: $\mathbf{x_1} = (1/2, 1/2, 0)$, $\mathbf{x_2} = \mathbf{e}_3$, and $\mathbf{x_3} = (2/7, 2/7, 3/7)$. The best response regions of these equilibria are depicted in Figure 1.2. The strategies to which $\mathbf{x_1}$ and $\mathbf{x_2}$, respectively, are best responses, are the regions below and above the horizontal line through $\mathbf{x_3}$. The two Nash equilibria $\mathbf{x_1}$ and $\mathbf{x_2}$ are superior to the points lying in their respective best response regions and are therefore evolutionarily stable. For the third, interior, equilibrium $\mathbf{x_3}$, however, no such neighbourhood exists, and hence it is not evolutionarily stable.

Can we modify this payoff matrix in a way such that this interior equilibrium becomes evolutionarily stable preserving evolutionary stability of the other two equilibria? Consider an ESS $\mathbf{x}$ and another strategy $\mathbf{y}$ with $\mathrm{Supp}(\mathbf{y}) \subseteq \mathrm{Supp}(\mathbf{x})$. Then, $\mathbf{y}$ is a best reply to $\mathbf{x}$ since it uses only pure strategies used by $\mathbf{x}$, and by Theorem 1.3, $\mathbf{x}^\top \mathbf{Ay} > \mathbf{y}^\top \mathbf{Ay}$ implying that $\mathbf{y}$ is not Nash.

Theorem 1.5. *If $\mathbf{x} \in \Delta^{ESS}$ and $\mathrm{Supp}(\mathbf{y}) \subseteq \mathrm{Supp}(\mathbf{x})$ for some $\mathbf{y} \neq \mathbf{x}$, then $\mathbf{y} \notin \Delta^{NE}$.*

This implies, that if some ESS is interior, it is unique. Also, since there is only a finite number of supports, Δ^{ESS} is finite.

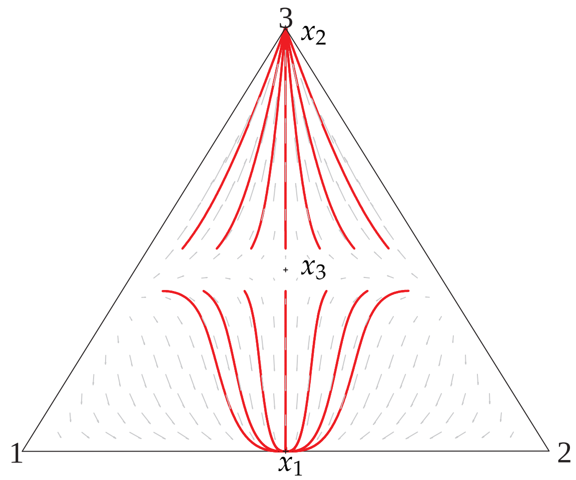

Figure 1.2: Solution orbits of Example 3.

Starting from the characterisation in terms of local superiority, we can generalise the definition of evolutionarily stable states to sets of strategies. We require every element of such a set to earn more against nearby strategies outside the set than these earn against themselves.

Definition 1.7. *A nonempty and closed set $X \subseteq \Delta$ is* evolutionarily stable *if for each $\mathbf{x} \in X$ there exists a neighbourhood U of $\mathbf{x}$ such that $\mathbf{x}^\top \mathbf{A}\mathbf{y} \geq \mathbf{y}^\top \mathbf{A}\mathbf{y}$ for all $\mathbf{y} \in U$ with strict inequality if $\mathbf{y} \notin X$.*

Observe that all elements of an evolutionarily stable set satisfy a weak local superiority condition in that the defining inequality $\mathbf{x}^\top \mathbf{A}\mathbf{y} \geq \mathbf{y}^\top \mathbf{A}\mathbf{y}$ is not strict as was required for local superiority. Replacing the inequalities in the characterisations of evolutionary stability above by weak inequalities, a respective set of characterisations can be shown. These strategies are then called *neutrally stable* and a set of characterisations exists, differing from those of evolutionary stability only in that inequalities are weak instead of strict.

1.2.3 Synthesis: Convergence Theorems

In the previous two sections we have seen two refinements of Nash equilibria: evolutionarily stable strategies and Nash equilibria that are asymptotically stable in the replicator dynamics. For the case of the generalised Rock-Scissors-Paper game $\mathbf{A}_{PRS}(\delta)$ we have seen that the first implies the second. We will now see that this is the case in general.

We cannot hope to find analytical solutions to the replicator dynamics given by Equation (1.1). Still there exists a powerful tool that will enable us to prove convergence properties without explicitly knowing the solution $\xi(\mathbf{x_0}, t)$. Consider a continuously differentiable vector field $v : D \mapsto \mathbb{R}^k$ with $D \subseteq \mathbb{R}^k$ and the associated dynamic system

$$\begin{aligned} \dot{\mathbf{x}} &= v(\mathbf{x}) \\ \mathbf{x}(0) &= \mathbf{x_0} \ . \end{aligned} \tag{1.4}$$

Again, we denote the solution for an initial value $\mathbf{x_0}$ by $\xi(\mathbf{x_0}, t)$. A natural way to prove that a fixed point is asymptotically stable in this system is by showing that the system looses "energy" over time. Suppose we are given such an energy function $E : D \mapsto \mathbb{R}$ which is minimised at a fixed point $\mathbf{x}$. If E decreases along any trajectory, any solution must eventually approach $\mathbf{x}$. In fact, asymptotically stable sets can be characterised in this way.

Theorem 1.6 (see, e. g., [14]). *Consider a dynamic system as given by Equation (1.4) and a closed set $C \subseteq D$. The set C is asymptotically stable if and only if there exists a neighbourhood U of C and a continuous function $E : D \mapsto \mathbb{R}_{\geq 0}$ such that $E(\mathbf{x}) = 0$ if and only if $\mathbf{x} \in C$, and*

$$E(\xi(\mathbf{x_0}, t)) < E(\mathbf{x_0}) \quad \text{if } \mathbf{x_0} \notin C,\ t > 0, \text{ and } \xi(\mathbf{x_0}, t') \in U \ \forall\, t' \in [0, t] \ .$$

The function E is called a strict *Lyapunov function*. Since we may not know the solution $\xi(\mathbf{x_0}, t)$ it may be hard to check whether $E(\xi(\mathbf{x_0}, t)) < E(\mathbf{x_0})$. Fortunately, it is often possible to check for this property without knowing $\xi(\cdot, \cdot)$ explicitly. The direction of a trajectory at $\mathbf{x}$ is given by the vector field $v(\mathbf{x})$ itself. To check whether $E(\mathbf{x})$ decreases in the direction $v(\mathbf{x})$ we make use of the gradient of E at $\mathbf{x}$,

$$\nabla E(\mathbf{x}) = \left(\frac{\partial E(\mathbf{x})}{\partial x_1}, \ldots, \frac{\partial E(\mathbf{x})}{\partial x_k} \right) \ .$$

If $v(\mathbf{x})$ points in the direction of $-\nabla E(\mathbf{x})$ this means that E decreases along the trajectory. If $v(\mathbf{x})$ is orthogonal to $\nabla E(\mathbf{x})$ the trajectory is tangential to a level curve of E and hence E does not change at all. In general, E decreases if and only if $v(\mathbf{x})$ makes an acute angle with $-\nabla E(\mathbf{x})$, or, equivalently, $\nabla E(\mathbf{x}) \cdot v(\mathbf{x}) < 0$. This yields the following theorem.

Theorem 1.7 (Lyapunov's direct method). *Consider a dynamic system as given by Equation (1.4) and a closed set $C \subseteq D$. If there exists a neighbourhood U of C and a continuously differentiable function $E : D \mapsto \mathbb{R}_{\geq 0}$ with $E(\mathbf{x}) = 0$ if and only if $\mathbf{x} \in C$ and*

$$\nabla E(\mathbf{x}) \cdot v(\mathbf{x}) < 0 \qquad \forall \mathbf{x} \notin C \ , \tag{1.5}$$

then the set C is asymptotically stable.

The requirement that the set C is closed is trivially met for the simple case that $C = \{\mathbf{x}\}$ is a singleton. Observe that

$$\nabla E(\mathbf{x}) \cdot v(\mathbf{x}) = \sum_{i=1}^{k} \frac{\partial E(x)}{\partial x_i} \cdot \frac{dx_i}{dt} = \dot{E}(\mathbf{x})$$

and hence condition (1.5) is equivalent to $\dot{E}(\mathbf{x}) < 0$.

In order to apply this technique to the replicator dynamics, we must find a suitable Lyapunov function. Fix a point $\mathbf{x}$ and consider the (Kullback-Leibler) relative-entropy function

$$H_{\mathbf{x}}(\mathbf{y}) = \sum_{i \in \mathrm{Supp}(\mathbf{x})} x_i \cdot \ln\left(\frac{x_i}{y_i}\right)$$

This function is defined for all $\mathbf{y}$ with $y_i \neq 0$ for all $i \in \mathrm{Supp}(\mathbf{x})$, i. e., the domain of $H_{\mathbf{x}}$ is a neighbourhood of $\mathbf{x}$ relative to Δ. First note that $H_{\mathbf{x}}(\mathbf{x}) = 0$. Now, let us substitute the replicator dynamics (1.1) into $H_{\mathbf{x}}$.

$$\begin{aligned} \dot{H}_{\mathbf{x}}(\mathbf{y}) &= -\sum_{i \in \mathrm{Supp}(\mathbf{x})} x_i \cdot \frac{1}{y_i} \cdot \dot{y}_i \\ &= -\lambda(\mathbf{y}) \cdot \sum_{i \in \mathrm{Supp}(\mathbf{x})} x_i \cdot \frac{1}{y_i} \cdot y_i \cdot ((\mathbf{A}\mathbf{y})_i - \mathbf{y}^\top \mathbf{A}\mathbf{y}) \\ &= \lambda(\mathbf{y}) \cdot (\mathbf{y}^\top \mathbf{A}\mathbf{y} - \mathbf{x}^\top \mathbf{A}\mathbf{y}) \ . \end{aligned}$$

We have seen that if $\mathbf{x}$ is evolutionarily stable, it is also locally superior (Theorem 1.4) implying that $\dot{H}_{\mathbf{x}}(\mathbf{y}) < 0$ for some neighbourhood U of $\mathbf{x}$ and $\mathbf{y} \in U$. The intersection of U with the domain of $H_{\mathbf{x}}$ is again a neighbourhood of $\mathbf{x}$ (relative to Δ). Furthermore, $H_{\mathbf{x}}$ is continuously differentiable. Altogether, $H_{\mathbf{x}}$ is a strict Lyapunov function.

Theorem 1.8 ([77]). *If $x \in \Delta^{ESS}$, $\mathbf{x}$ is asymptotically stable in the replicator dynamics (1.1).*

We see that evolutionarily stable strategies attract nearby strategies. Consider an interior ESS $\mathbf{x} \in \Delta^{ESS} \cap \mathrm{int}(\Delta)$. For this strategy, all strategies are best replies, and $H_{\mathbf{x}}(\mathbf{y})$, which is defined on the entire interior, is negative for all $\mathbf{y} \in \mathrm{int}(\Delta)$. Consequently, the vector field of the replicator dynamics points strictly into every contour set and hence, $\mathbf{x}$ attracts the entire interior.

Theorem 1.9. *If $\mathbf{x} \in \Delta^{ESS} \cap \mathrm{int}(\Delta)$, then $\xi(\mathbf{x_0}, t)_{t \to \infty} \to \mathbf{x}$ for any $\mathbf{x_0} \in \mathrm{int}(\Delta)$.*

The proof of Theorem 1.8 can be generalised to evolutionarily stable sets X by defining a new Lyapunov function as the minimum of all $H_{\mathbf{x}}(\mathbf{y})$

over all $\mathbf{x}$ where it is defined:

$$H(\mathbf{y}) = \min_{\mathbf{x}\in X:\mathrm{Supp}(\mathbf{x})\subseteq\mathrm{Supp}(\mathbf{y})} H_{\mathbf{x}}(\mathbf{y}) \ .$$

Using this Lyapunov function it can be shown that evolutionarily stable sets $S \subseteq \Delta$ are asymptotically stable.

In the literature, also discrete versions of the replicator dynamics are discussed. For these dynamics, however, only weaker convergence properties hold due to overshooting effects. E. g., it is not even guaranteed that dominated strategies vanish, as is shown by the Dekel-Scotchmer Example [23]. We will consider similar effects in Chapter 3.

1.2.4 Related Work

From the perspective of theoretical computer science we are interested in the complexity of computing evolutionarily stable strategies. By a reduction from CLIQUE it can be shown that finding an evolutionarily stable strategy for a given symmetric two player game is both NP-hard and co-NP-hard [26, 55]. It is even co-NP-hard to determine whether any given strategy is an evolutionarily stable strategy. In [26] it is furthermore shown, that counting the number of evolutionarily stable strategies is #P-hard.

1.3 Selfish Routing

In 1952, Wardrop [83] (for a good introduction see also [69]) proposed a model of selfish traffic in networks. Whereas it was originally introduced for the study of road traffic, it is also well-suited for the analysis of digital traffic in communication networks. This model will be the foundation of all subsequent analyses and modelling. We present it here, along with its equilibrium concept and a useful potential function characterisation which we will make frequent use of in the main part of this thesis.

1.3.1 The Wardrop Model

An *instance* of the Wardrop routing game is given by a triple $\Gamma = (G, \mathcal{C}, (\ell_e))$ composed of a finite directed multi-graph $G = (V, E)$, a family of differentiable, non-decreasing, non-negative *latency functions* $(\ell_e)_{e\in E}$, $\ell_e : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$, and a family of k *commodities* $\mathcal{C} = (C_i)_{i\in[k]}$. Any commodity $C_i = (s_i, t_i, r_i)$ specifies a source node $s_i \in V$, a target node $t_i \in V$, and a flow demand $r_i \in \mathbb{R}_+$ that is to be routed from s_i to t_i. The total flow demand is denoted by $r = \sum_{i\in[k]} r_i$. For $i \in [k]$, let $\mathcal{P}_i$ denote the set of acyclic paths in G that start at s_i and end at t_i, and let $\mathcal{P} = \bigcup_{i\in[k]} \mathcal{P}_i$. We may assume that the source-sink pairs are pairwise disjoint and, hence, so are the $\mathcal{P}_i$. Then, each path $P \in \mathcal{P}$ uniquely determines a commodity i_P such that $P \in \mathcal{P}_{i_P}$. An instance is *symmetric* or *single-commodity* if $k = 1$ and *asymmetric* or *multi-commodity* if $k > 1$. The maximum path length is denoted by $D = \max_{P\in\mathcal{P}} |P|$. We say that Γ is a *singleton instance* if $D = 1$.

A flow vector $(f_P)_{P\in\mathcal{P}}$ is *feasible* for an instance Γ if it is non-negative and satisfies the flow demands, i. e.,

$$\begin{aligned} \sum_{P\in\mathcal{P}_i} f_P &= r_i \qquad \forall i \in [k] \\ f_P &\geq 0 \qquad \forall P \in \mathcal{P} \ . \end{aligned} \tag{1.6}$$

Let $\mathcal{F}_\Gamma$ denote the polyhedron specified by these constraints. The path flow vector f uniquely determines edge flows

$$f_e = \sum_{P\ni e} f_P \qquad \text{and} \qquad f_{e,i} = \sum_{P\in\mathcal{P}_i : e\in P} f_P$$

for all edges $e \in E$ and all commodities $i \in [k]$.

A flow vector f induces latencies on the edges. The latency of edge $e \in E$ is $\ell_e(f) = \ell_e(f_e)$. Latency is additive along a path, i. e., for any path $P \in \mathcal{P}$, $\ell_P(f) = \sum_{e\in P} \ell_e(f)$. For each commodity $i \in [k]$ this implies an *average latency* of

$$L_i(f) = \sum_{P\in\mathcal{P}_i} \frac{f_P}{r_i} \cdot \ell_P(f) = \sum_{e\in E} \frac{f_{e,i}}{r_i} \cdot \ell_e(f) \ .$$

The average latency of the entire flow is then

$$L(f) = \sum_{P\in\mathcal{P}} \frac{f_P}{r} \cdot \ell_P(f) = \sum_{e\in E} \frac{f_e}{r} \cdot \ell_e(f) = \sum_{i\in[k]} \frac{r_i}{r} \cdot L_i(f) \ .$$

Some results presented in the main part of this thesis depend on extremal values of the latency functions. We define these quantities here, though we will repeat the definitions when needed. The optimal average latency for an instance is given by $L_{\min} = \min_{f\in\mathcal{F}_\Gamma} L(f)$, and the maximum average latency is given by $L_{\max} = \max_{f\in\mathcal{F}_\Gamma} L(f)$. The optimal average latency of a single commodity is given by $\tilde{L}_{\min} = \min_{i\in[k]} \min_{f\in\mathcal{F}_\Gamma} L_i(f)$. Finally, the maximum latency that can be induced on a path is given by $\ell_{\max} = \max_{P\in\mathcal{P}} \max_{f\in\mathcal{F}_\Gamma} \ell_P(f)$.

Let us remark that the combinatorial structure of this model can be generalised easily in the following way. Instead of considering a set E of edges in a network we can simply assume that E is a set of arbitrary *resources*.

Then, any collection of resources is a possible *strategy*, and the strategy space of an agent can be an arbitrary subset $\mathcal{P}_i \subseteq 2^E$. In this case, we can specify an instance by the triple $\Gamma = (E, \mathcal{C}, (\ell_e)_{e \in E})$ where any commodity $C_i = (\mathcal{P}_i, r_i)$ specifies the strategy set $\mathcal{P}_i$ explicitly rather than implicitly in terms of s_i and t_i. We will use this generalisation a few times in Chapter 4. Our analyses are independent of the combinatorial structure of the strategy space and do not rely on the assumption that it is given by a graph. However, if we view an instance Γ as an input of an algorithm, the specification of the strategy spaces by sources and sinks in a graph yields a compact way to represent an exponentially large strategy space. Throughout this thesis, we use the terms edges ($e \in E$) and paths ($P \in \mathcal{P}$) rather than the terms resources and strategies.

1.3.2 Wardrop Equilibria

In this setting, two possible objectives come to mind. The classical objective would be to optimise for the total benefit or *social welfare* of participants in the network, e. g., to minimise the total incurred latency:

$$\min_{f \in \mathcal{F}_\Gamma} L(f) \quad . \tag{1.7}$$

The feasible set $\mathcal{F}_\Gamma$ is convex, and the objective function L is convex if the functions $x \cdot \ell_e(x)$ are convex. Thus, we can compute minima in polynomial time, e. g., using the ellipsoid method [46, 74]. Note that the specification of the feasible set as specified by Equation (1.6) contains an exponential number of constraints. This can easily be reduced to a polynomial number by formulating it in terms of the edge flows and specifying the flow conservation constraints for each commodity.

Taking the game theoretic perspective we are not interested in minimising the social welfare, but we are rather interested in finding a balanced, or fair, allocation. For a finite number of agents competing for resources, the latency functions can be described by an N-dimensional payoff table where N is the number of users. Thus, we can model this scenario as a strategic game. The Wardrop model is the limit of such a game as N approaches infinity. We envision the flow as composed of an infinite number of agents each of which carries an infinitesimal amount of flow. Generalising the definition of Nash equilibria to games with an infinite number of players, we cannot require that no *single* player can unilaterally improve their situation by changing their strategy. Instead we require that no arbitrarily small portion of flow can be shifted from its path to another without increasing its latency. Formally, for any flow vector f let $\tilde{f}^{P_1,P_2,\epsilon}$ denote such a flow where a flow amount of ϵ is shifted from path P_1 to path P_2, i. e.,

$$\tilde{f}_P^{P_1,P_2,\epsilon} = \begin{cases} f_P - \epsilon & \text{if } P = P_1 \\ f_P + \epsilon & \text{if } P = P_2 \\ f_P & \text{otherwise.} \end{cases}$$

Definition 1.8 (Nash equilibrium). *A flow vector $(f_P)_{P \in \mathcal{P}}$ is at a* Nash equilibrium *if for every commodity $i \in [k]$ and all paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that for all $\epsilon \in [0, f_{P_1}]$, $\ell_{P_1}(f) \leq \ell_{P_2}(\tilde{f}^{P_1,P_2,\epsilon})$. The set of Nash equilibria in the Wardrop model is denoted $\mathcal{F}_\Gamma^{NE}$.*

Letting ϵ tend to zero we obtain a characterisation of Nash equilibria in the Wardrop model that is more comfortable.

Theorem 1.10 (Wardrop equilibrium [40]). *A flow vector f is at a Nash equilibrium if for all commodities $i \in [k]$ and all $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f) \leq \ell_{P_2}(f)$.*

This characterisation is also referred to as a *Wardrop equilibrium* or as *Wardrop's principle* [40].

From the definition it follows that at a Wardrop equilibrium each path P is either unused or it has minimal, and therefore also average, latency L_{i_P}. This fact discloses a connection between Wardrop equilibria and flows optimising $L(f)$. A flow is (locally and globally) optimal if shifting flow from one path to another can only increase L. In other words, on all used paths the *marginal cost* specifying the infinitesimal contribution that additional flow causes on such a path, is minimal. This similarity between the characterisation of Nash and optimal flows is formalised in the following theorem. For any function $\ell(x)$ let $\ell^*(x) = (x \cdot \ell(x))' = \ell(x) + x \cdot \ell'(x)$ denote its marginal cost function.

Theorem 1.11 ([83, 70]). *Consider an instance $\Gamma = (G, (\ell_e)_{e \in E}, \mathcal{C})$ such that for all $e \in E$, $\ell_e(x) \cdot x$ is convex. Then, a feasible flow f is optimal for Γ if and only if it is at a Wardrop equilibrium for the instance $(G, (\ell_e^*)_{e \in E}, \mathcal{C})$ with marginal cost latency functions.*

1.3.3 The Beckmann-McGuire-Winsten Potential

We can reverse Theorem 1.11 to obtain a characterisation of Wardrop equilibria in terms of a potential function due to Beckmann, McGuire, and Winsten [11]. Consider the function $h_e(x) = 1/x \cdot \int_0^x \ell_e(u)\, du$ and observe that the marginal cost function of h_e is the function ℓ_e again. The total cost with respect to the functions $(h_e)_{e \in E}$ is

$$\Phi(f) = \sum_{e \in E} f_e \cdot h_e(f_e) = \sum_{e \in E} \int_0^{f_e} \ell_e(u)\, du \quad . \tag{1.8}$$

Consequently, a flow is at a Wardrop equilibrium with respect to $(\ell_e)_{e\in E}$ if and only if it minimises Φ, i. e., it solves the convex program

$$\min_{f\in\mathcal{F}_\Gamma} \Phi(f) \ . \tag{1.9}$$

We denote the minimum potential by $\Phi^* = \min_{f\in\mathcal{F}_\Gamma} \Phi(f)$. Since the feasible set is convex and non-empty, we obtain:

Theorem 1.12. *For any instance $\Gamma = (G, (\ell_e)_{e\in E}, \mathcal{C})$ there exists a flow f at Wardrop equilibrium. Furthermore for any two Wardrop equilibria f and f' and every edge $e \in E$, $\ell_e(f) = \ell_e(f')$.*

The second statement follows from the fact that $f_e \cdot h_e$ must be linear in between f_e and f'_e since otherwise a convex combination of f and f' would achieve a smaller objective function value. Consequently, ℓ_e is constant within this interval.

The potential Φ also has an intuitive interpretation which will guide many of our analyses in the main part. Consider any flow f and assume that an infinitesimal flow dx is inserted on path P. This increases the potential by $\ell_P(f)\,dx$. A similar statement holds for removing flow. This observation is useful for computing the potential gain that is achieved by shifting flow from one path to another. If a flow of dx is moved from path P to path Q, the potential changes by $(\ell_Q - \ell_P)\,dx$.

1.3.4 Related Work

A natural question concerning the Wardrop model asks for the degradation of performance due to the selfishness of the agents. To that end, one considers the *price of anarchy* defined as the ratio between the total latency of a worst Wardrop equilibrium and the minimum total latency

$$\rho(\Gamma) = \frac{\max_{f\in\mathcal{F}_\Gamma^{NE}} L(f)}{\min_{f\in\mathcal{F}_\Gamma} L(f)} \ .$$

For affine latency functions, Roughgarden and Tardos [70] show that $\rho(\Gamma) \le 4/3$. For polynomial latency functions with degree d, $\rho(\Gamma) \le \mathcal{O}\left(d/(\log d)\right)$. These upper bounds are tight and instances with two nodes and two edges exist that achieve these values [67]. Whereas in [70] and [67] it is assumed that the strategy space is the set of paths in a network, many of the results also hold in the general case [71].

One can consider the price of anarchy also with respect to objective functions other than the average latency L. In [68] it is shown that under the objective of minimising the maximum latency rather than the average, the price of anarchy is at most $|V| - 1$ for single-commodity instances.

Cominetti *et al.* [21] consider network games with users which control a non-negligible amount of traffic which may be split among several paths. For these games, the price of anarchy is 3/2 for affine latency functions. For a finite number of users that are not allowed to split their flow, bounds on the price of anarchy are presented in [6, 19]. In [41] the authors consider games in which groups of users may collaborate and introduce the price of collusion which measures the degradation of performance due to this behaviour.

A notion converse to the price of anarchy is discussed in [66]. Therein, bounds on the unfairness of optimal routing, measured as the maximal ratio between a path latency at optimum and at Nash equilibrium, are given.

Knowing that that price of anarchy may be large, a natural question is how this ratio can be decreased by external intervention. Several approaches have been proposed. In [20] and [35] a scenario is considered in which taxes are imposed on the users such that Nash equilibria coincide with Wardrop equilibria, and bounds on the magnitude of such taxes are presented. Another possibility to decrease the price of anarchy is to control a fraction of the individuals centrally. The tradeoff between the amount of centrally controlled individuals and the price of anarchy is studied in [65].

Rosenthal [63] uses a discrete version of the potential function Φ to show existence of pure Nash equilibria in discrete congestion games. Fabrikant *et al.* [29] consider the complexity of computing Nash equilibria in a discrete model. They show that computing Nash equilibria is PLS-complete in general whereas there exists a polynomial time algorithm for the case of symmetric network congestion games. The first is reproved in [1], which also gives a nice characterisation of congestion games that allow convergence towards Nash equilibria by a polynomial length sequence of best response moves.

1.4 Outline of the Thesis

As has been discussed above, we believe that equilibria in the Wardrop model are meaningful only if they are motivated by a simple operational model of the agents' behaviour. This model should rely on assumptions that are significantly milder than the ones needed for the motivation of Nash equilibria for general strategic games. As a first step we apply evolutionary population dynamics to the Wardrop model. In Chapter 2 we see that Wardrop equilibria are evolutionarily stable. In the single-commodity case we can thus show that a population following the replicator dynamics actually approaches the set of Wardrop equilibria. In the multi-commodity case, we generalise the replicator dynamics accordingly, taking an approach that slightly differs from the standard approach of evolutionary game theory for asymmetric games. Thus we are able to prove convergence towards Wardrop equilibria also in the multi-commodity case. In addition, we present bounds on the time to reach approximate Wardrop equilibria. The continuous-time model of the replicator dynamics idealistically assumes that the effect of the agents' decisions can be observed immediately. Since this is not the case in computer communication networks, the time bounds presented in this chapter must be considered preliminary.

Subsequently, we consider a model in which information is observed with a delay [52] to resolve this shortcoming. It is well-known in both practise and theory that decisions based on stale latency information may cause oscillation effects and harm network performance seriously. In Chapter 3 we consider a fairly general class of dynamics. We identify sufficient properties of the dynamics and the latency functions that guarantee convergence towards Wardrop equilibria. It is not surprising that this condition depends on the steepness of the latency functions: The more sensitive the latency functions are to minor shifts of flow, the more vulnerable the dynamics becomes to overshooting. We will see that if our dynamics is slowed down by a factor that is reciprocal to the maximum slope of the latency functions, it converges towards the set of Wardrop equilibria. Thus, any bound on the time to reach approximate equilibria based on this criterion has a pseudopolynomial flavour.

In Chapter 4 we introduce a particular rerouting policy which applies a clever sampling technique in order to avoid the dependence on the maximum slope. This policy approaches approximate Wardrop equilibria much quicker than policies from the large class considered in Chapter 3 do. To measure the convergence time, we study two kinds of approximate equilibria. First, we consider allocations in which all but a small fraction of agents deviate from the average latency of their commodity by no more than a small fraction of the overall average latency. Remarkably, the time to reach bicriterial approximate equilibria of this kind is largely independent of the network size. However, these approximate equilibria are transient. Approximate equilibria that are persistent with respect to our dynamics can be defined in terms of the Beckmann-McGuire-Winsten potential. For symmetric instances, we show that our dynamics reaches a flow which is within $(1+\epsilon)$ of the optimal potential in polynomial time. The time bounds presented in this chapter show that the critical parameter is not the maximum slope of the latency functions, but rather their elasticity.

In the final chapter of this thesis we exchange the theorist's perspective for a more applied one. In Chapter 5 we present a traffic engineering scheme, REPLEX, inspired by the population dynamics discussed so far. This protocol was implemented and evaluated in a simulation framework with realistic traffic properties including a full TCP implementation at packet level and HTTP workload that is generated on the basis of realistic statistical properties. Simulation results on both artificial and real network topologies indicate that the protocol converges quickly and significantly improves network performance.

1.4.1 Bibliographical Notes

Many of the results presented in the main part of this thesis have been presented as joint work in preliminary form at various conferences. Results about evolutionary stability of and convergence towards Wardrop equilibria (Chapter 2) have been presented in [33]. A convergence analyses in a model of stale information (Chapter 3) was presented in [34]. The results of Chapter 4 regarding the speed of convergence have been sketched in [32]. Our traffic engineering protocol REPLEX (Chapter 5) has been introduced in [31].

CHAPTER 2

Evolution of Selfish Routing

In the previous chapter we have seen how the notion of Wardrop equilibria captures the effect of selfishness in communication networks that lack central coordination. Though Wardrop equilibria are stable once they are reached, the assumptions made to attain them cannot be justified for real-world applications, the prime example being the Internet. Recall these assumptions:

1. *Complete and accurate information about the payoff functions.* In our scenario, payoff corresponds to (inverse) latency. Latency functions, however, are not known to the agents in real-world networks. Neither are the necessary parameters of the links publicly available, nor can they be measured easily by the agents. Furthermore, agents don't have *perfect information*, since they cannot directly observe the actions of the other agents.

2. *Common knowledge of unbounded rationality.* Agents cannot make reliable assumptions on the rationality of other agents. Other agents may be rational, but they may also act faulty, randomly, or even maliciously.

3. *Unlimited reasoning capabilities.* Computational power is a scarce resource, in particular for Internet routers which are responsible for the path selection. Even if agents had the necessary information, we cannot expect them to compute a Wardrop equilibrium off-line and then route their flow according to the result of this computation.

To avoid these unrealistic assumptions we apply techniques from evolutionary game theory to the Wardrop model. We start by modelling a dynamic population of selfish agents executing simple adaption rules resulting in a dynamical system similar to the replicator dynamics. These rules rely on very mild assumptions only, which we will discuss below. Subsequently, we will derive stability and convergence properties of this dynamics.

2.1 The Replicator Dynamics in the Wardrop Model

In this section we show how the replicator dynamics as presented in the previous section translates naturally into the model of selfish routing. We start with the single-commodity case, i. e., $k = 1$. We consider an infinite population of agents each of which chooses a particular path $P \in \mathcal{P}$. The fraction of agents choosing path P is denoted by f_P. Since all agents carry the same infinitesimal amount of flow, the vector $(f_P)_{P \in \mathcal{P}}$ can be interpreted both as a population vector and as a flow vector induced by this population.

We devise a rerouting policy (and subsequently, a differential equation) that specifies the behaviour of f over time. We assume that agents become active at certain Poisson rates. Consider an agent that is activated and currently uses path $P \in \mathcal{P}$. It performs two steps:

1. *Sampling:* Sample another agent uniformly at random. This agent uses path $Q \in \mathcal{P}$ with probability f_Q/r.
2. *Migration:* Migrate from path P to path Q if $\ell_P(f) > \ell_Q(f)$ with probability proportional to the latency improvement $\ell_P(f) - \ell_Q(f)$.

Since the total volume of agents utilising path $P \in \mathcal{P}$ is f_P and all of these agents act according to the above rules, the total rate at which agents migrate from path P to Q is

$$\lambda(f) \cdot \frac{f_P \cdot f_Q}{r} \cdot (\ell_P(f) - \ell_Q(f)) \ ,$$

where the scalar function $\lambda(f)$ accounts for the constant of proportionality in the migration probability (step 2) and the Poisson rate at which agents are activated. Note that this expression holds for both $\ell_P(f) > \ell_Q(f)$ and $\ell_P(f) < \ell_Q(f)$. Hence, the rate at which the population share utilising path P changes is

$$\dot{f}_P = \lambda(f) \cdot f_P \sum_{Q \in \mathcal{P}} \frac{f_Q}{r} \cdot (\ell_Q(f) - \ell_P(f)) \ .$$

Using the definition of the average latency $L(f)$ and $\sum_{Q \in \mathcal{P}} f_Q = r$, we obtain the following characterisation of our dynamic system:

$$\begin{aligned} \dot{f}_P &= \lambda(f) \cdot f_P \cdot (L(f) - \ell_P(f)) \qquad \text{for all } P \in \mathcal{P} \qquad (2.1) \\ f(0) &= f_0 \ . \end{aligned}$$

Comparing Equations (1.1) and (2.1), we see that this is a variant of the replicator dynamics where latency corresponds to negative payoff. The expected payoff, however, is not linear in f unless all ℓ_e have the form $\ell_e(x) = a_e \cdot x$ for some $a_e \in \mathbb{R}$. Let us remark that our simple rerouting policy presented above is, again, only one out of many possible motivations for this dynamics. The imitation rules described in Section 1.2.1 can also be applied to selfish routing and lead to the same kind of dynamics differing only in the scaling factor $\lambda(f)$.

The choice of λ does not have any influence on the solution orbit as long as it is Lipschitz continuous and strictly positive. It merely influences the speed at which the orbit is traversed, and, consequently, the time of convergence, which we will discuss in Section 2.3.

Multi-commodity networks correspond to asymmetric games. Generalisations of the replicator dynamics for asymmetric games proposed in the literature do not work well with the Wardrop model. In general strategic games, the expected payoff of a pure strategy for a player of a given type depends on the population mixture of the subpopulations of other player types, but not on the population mixture of its own type. In the Wardrop model, where player types correspond to commodities, this is different. Here, the latency of a player does also depend on the routing decisions made by other players of the same commodity.

Despite of this, Equation (2.1) can be generalised to multi-commodity networks. Consider an agent of commodity $i \in [k]$ currently using path $P \in \mathcal{P}_i$. If we apply the same rules as above, the probability to sample path Q is f_Q/r_i. Again, using the definition of the average latency $L_i(f)$ and $\sum_{Q \in \mathcal{P}_i} f_Q = r_i$, we get:

$$\begin{aligned} \dot{f}_P &= \lambda_i(f) \cdot f_P \cdot (L_i(f) - \ell_P(f)) \qquad \text{for all } P \in \mathcal{P} \\ f(0) &= f_0 \ , \end{aligned} \tag{2.2}$$

where we use a different scaling factor $\lambda_i(f)$ for every commodity.

We believe that this is a reasonable model of communication and selfish behaviour for several reasons.

1. *Simplicity.* The reasoning involved in this process is limited to few and simple comparison and sampling steps.

2. *Locality.* The necessary information, namely the flow and latency values of paths of the same commodity, is available at the source node. Agents need to communicate with agents sharing the same source and destination nodes only.

3. *Selfish steps.* The steps made by the agents are selfish and myopic in that they exchange a path only for a better one.

4. *Statelessness.* The process does not require the agents to maintain any state information. The agents' decisions depend merely on the current situation and not on history.

This policy does not depend on accurate knowledge of the network and its latency functions as well as unbounded rationality and reasoning capabilities. Note, however, that our population model still assumes perfect information in that actions made by other agents become visible to the other agents immediately. In subsequent chapters we will see how we can abandon this assumption as well.

2.1.1 Preliminaries

We can easily show that with respect to the dynamic system given by Equation (2.2), the feasible set $\mathcal{F}_\Gamma$ is invariant. Consider a commodity $i \in [k]$. We show that $\sum_{P\in\mathcal{P}_i} \dot{f}_P = 0$ implying $\sum_{P\in\mathcal{P}_i} f_P = r_i$ if $f_0 \in \mathcal{F}_\Gamma$.

$$\begin{aligned}
\sum_{P\in\mathcal{P}_i} \dot{f}_P &= \sum_{P\in\mathcal{P}_i} \lambda_i(f) \cdot f_P \cdot (L_i(f) - \ell_P(f)) \\
&= \lambda_i(f) \cdot \left(r_i \cdot L_i(f) - r_i \sum_{P\in\mathcal{P}_i} \frac{f_P}{r_i} \cdot \ell_P(f) \right) \\
&= \lambda_i(f) \cdot (r_i \cdot L_i(f) - r_i \cdot L_i(f)) \\
&= 0 \ .
\end{aligned}$$

Furthermore, clearly $f_P(t) \geq 0$ for all $t \geq 0$. Hence $f(t) \in \mathcal{F}_\Gamma$ for all $t \geq 0$ if $f_0 \in \mathcal{F}_\Gamma$.

We adapt the definition of evolutionary stability to the Wardrop model. As we have seen, Wardrop equilibria are not necessarily unique. Due to our characterisation in terms of a convex program, we know that the set of Wardrop equilibria is closed. We use a definition similar to our setwise generalisation of evolutionary stability from Section 1.2.2 which is based on local superiority.

Definition 2.1 (evolutionarily stable)**.** *For an instance Γ, a nonempty and closed set $X \subseteq \mathcal{F}_\Gamma$ is* evolutionarily stable *if for all $f \in X$ there exists a neighbourhood U of f such that $f \cdot \ell(f') < f' \cdot \ell(f')$ for all $f' \in U \setminus X$ and $f \cdot \ell(f') \leq f' \cdot \ell(f')$ for $f' \in X$.*

We will see that this adaption of the notion of evolutionary stability is suitable for our purposes as it implies convergence of the replicator dynamics to the set of Wardrop equilibria in symmetric instances of the Wardrop routing game.

2.2 Evolutionary Stability and Convergence

In this section, we study stability and robustness aspects of Wardrop equilibria. We prove that the set of Wardrop equilibria is evolutionarily stable. Based on this property, we can then prove that solutions of the replicator dynamics converge towards Wardrop equilibria.

2.2.1 Evolutionary Stability

Lemma 2.1. *The set of Wardrop equilibria $\mathcal{F}_\Gamma^{NE}$ is evolutionarily stable. The neighbourhood U in Definition 2.1 can be chosen to be $U = \mathcal{F}_\Gamma$.*

Proof. Fix a Wardrop equilibrium $f \in \mathcal{F}_\Gamma^{NE}$ for instance Γ. We show that for any population $\tilde{f} \in \mathcal{F}_\Gamma$, it holds that $f \cdot \ell(\tilde{f}) \leq \tilde{f} \cdot \ell(\tilde{f})$. Furthermore, if $\tilde{f} \notin \mathcal{F}_\Gamma^{NE}$, the inequality is strict.

Fix any commodity $i \in [k]$. By definition of Wardrop equilibria, the latency induced by f on all paths in $\mathcal{P}_i$ may be equal or greater, but not less than the latency of any used path in $\mathcal{P}_i$. Hence, for every population $\tilde{f}$,

$$
\begin{aligned}
f \cdot \ell(f) &= \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i} f_P \, \ell_P(f) \\
&= \sum_{i \in [k]} r_i \, L_i(f) \\
&= \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i} \tilde{f}_P \, L_i(f) \\
&\leq \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i} \tilde{f}_P \, \ell_P(f) \\
&= \tilde{f} \cdot \ell(f) \ .
\end{aligned}
$$

The inequality holds because for any path $P \in \mathcal{P}_i$ we have $\ell_P(f) \geq L_i(f)$ as discussed above. As a consequence,

$$
\begin{aligned}
\tilde{f} \cdot \ell(\tilde{f}) &\geq f \cdot \ell(f) + \tilde{f} \cdot \ell(\tilde{f}) - \tilde{f} \cdot \ell(f) \\
&= f \cdot \ell(f) + \sum_{P \in \mathcal{P}} \tilde{f}_P (\ell_P(\tilde{f}) - \ell_P(f)) \\
&= f \cdot \ell(f) + \sum_{e \in E} \tilde{f}_e (\ell_e(\tilde{f}) - \ell_e(f)) \ . \qquad (2.3)
\end{aligned}
$$

Consider an edge $e \in E$. There are three cases:

1. $\tilde{f}_e > f_e$. Because of monotonicity of ℓ_e, it holds that $\ell_e(\tilde{f}_e) \geq \ell_e(f_e)$. Therefore also $\tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) \geq f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f))$.

2. $\tilde{f}_e < f_e$. Because of monotonicity of ℓ_e, it holds that $\ell_e(\tilde{f}_e) \leq \ell_e(f_e)$. Again, $\tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) \geq f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f))$.

3. $\tilde{f}_e = f_e$. Then, $\tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) = f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f))$.

In all cases we have $\tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) \geq f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f))$. Altogether, starting from Equation (2.3),

$$\begin{aligned}
\tilde{f} \cdot \ell(\tilde{f}) &\geq f \cdot \ell(f) + \sum_{e \in E} \tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) \\
&\geq f \cdot \ell(f) + \sum_{e \in E} f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) \\
&= f \cdot \ell(f) + f \cdot \ell(\tilde{f}) - f \cdot \ell(f) \ , \\
&= f \cdot \ell(\tilde{f}) \ ,
\end{aligned}$$

which is our claim.

It remains to show that the inequality is strict if $\tilde{f}$ is not a Wardrop equilibrium. There are two cases.

1. $\ell(f) = \ell(\tilde{f})$, i. e., for all $P \in \mathcal{P}$, $\ell_P(f) = \ell_P(\tilde{f})$. Since within any commodity $i \in [k]$ all paths used by f have the same latency $L_i(f)$, and since $\tilde{f}$ is not a Wardrop equilibrium, there must exist a commodity $i \in [k]$ and path $P \in \mathcal{P}_i$ with $\tilde{f}_P > 0$ and $\ell_P(\tilde{f}) = \ell_P(f) > L_i(f)$. Since there does not exist a path P with $\ell_P(f) < L_i(f)$, this implies $f \cdot \ell(\tilde{f}) < \tilde{f} \cdot \ell(\tilde{f})$.

2. $\ell(f) \neq \ell(\tilde{f})$, i. e., there exists a path $P \in \mathcal{P}$ such that $\ell_P(f) \neq \ell_P(\tilde{f})$. Then there exists an edge $e \in E$ with $\ell_e(f_e) \neq \ell_e(\tilde{f}_e)$ and consequently $f_e \neq \tilde{f}_e$. Then, we are in case 1 or 2 of the above case differentiation for this edge, and the respective inequality is strict, i. e., $\tilde{f}_e \cdot (\ell_e(\tilde{f}) - \ell_e(f)) > f_e \cdot (\ell_e(\tilde{f}) - \ell_e(f))$.

In both cases, $f \cdot \ell(\tilde{f}) < \tilde{f} \cdot \ell(\tilde{f})$. □

2.2.2 Convergence for Single-Commodity Instances

Using Lyapunov's direct method in conjunction with Lemma 2.1 we can show that solutions of the replicator dynamics converge towards a Wardrop equilibrium if the initial population f_0 lies in the interior of the simplex, i. e., there are no unused strategies. As a Lyapunov function we use the well-known (Kullback-Leibler) relative entropy measure.

Theorem 2.2. *Consider a single-commodity instance Γ. Let λ be a continuous, positive function. Let $\xi(f_0, t)$ be a solution to the replicator dynamics (2.1) for some initial population $f_0 \in \text{int}(\mathcal{F}_\Gamma)$. Then, $\xi(f_0, t)_{t \to \infty} \to \mathcal{F}_\Gamma^{NE}$.*

Proof. We use the relative-entropy measure $H_f : \mathrm{int}(\mathcal{F}_\Gamma) \to \mathbb{R}_{\geq 0}$ as a Lyapunov function:

$$H_f(\xi) = \sum_{P \in \mathrm{Supp}(f)} f_P \cdot \ln \frac{f_P}{\xi_P} \ .$$

Note that $H_f(\xi)$ is well-defined for all $\xi \in \mathrm{int}(\mathcal{F}_\Gamma)$ since $\xi_P > 0$ and $f_P > 0$ for all $P \in \mathrm{Supp}(f)$. Since we can write $H_f(\xi) = \sum_{P \in \mathrm{Supp}(f)} (f_P \cdot \ln(f_P) - f_P \cdot \ln(\xi_P))$, we have

$$\dot{H}_f(\xi) = -\sum_{P \in \mathrm{Supp}(f)} f_P \cdot \frac{\dot{\xi}_P}{\xi_P} \ .$$

Substituting the replicator dynamics (2.2) for $\dot{\xi}_P$, we get

$$\begin{aligned}
\dot{H}_f(\xi) &= \lambda(\xi) \sum_{P \in \mathrm{Supp}(f)} f_P \left(\ell_P(\xi) - L(\xi)\right) \\
&= \lambda(\xi) \left(\sum_{P \in \mathrm{Supp}(f)} f_P \, \ell_P(\xi) - r\, L(\xi) \right) \\
&= \lambda(\xi) \left(\sum_{P \in \mathrm{Supp}(f)} f_P \, \ell_P(\xi) - \sum_{P \in \mathcal{P}} \xi_P \, \ell_P(\xi) \right) \\
&= \lambda(\xi) \sum_{P \in \mathrm{Supp}(f)} (f_P - \xi_P) \cdot \ell_P(\xi) \\
&= \lambda(\xi) \cdot (f - \xi) \cdot \ell(\xi) \ . \qquad (2.4)
\end{aligned}$$

By Lemma 2.1, every term of this equation is non-positive and is zero only if ξ is a Wardrop equilibrium.

If the Wardrop equilibrium is unique, the proof of convergence is a direct application of Lyapunov's direct method (Theorem 1.7) with $C = \{f\}$ where f is the unique Wardrop equilibrium. Since Wardrop equilibria are not necessarily unique in our case, we consider the function

$$\mathcal{H}(\xi) = \min_{f \in \mathcal{F}_\Gamma^{NE}} H_f(\xi) \ .$$

Note that $\mathcal{H}(\xi)$ is continuous in ξ and $\mathcal{H}(\xi) \geq 0$ with equality only for $\xi \in \mathcal{F}_\Gamma^{NE}$. From the above, we know that for all $f_0 \in \mathrm{int}(\mathcal{F}_\Gamma)$, $H(\xi(f_0, t)) < H(f_0)$ for all $t > 0$. Hence, $\mathcal{H}$ and $\xi(f_0, \cdot)$ satisfy the conditions of Theorem 1.6. Since $\xi(f_0, t) \in \mathrm{int}(\mathcal{F}_\Gamma)$ for all $t \geq 0$, the vector field of the replicator dynamics (2.1) points strictly into every contour set of $\mathcal{H}$ and hence $\xi(f_0, t)_{t \to \infty} \to \mathcal{F}_\Gamma^{NE}$. □

2.2.3 Convergence of Multi-Commodity Instances

Note that Lemma 2.1 immediately covers the multi-commodity case. Despite of this, this does not suffice to prove convergence in the multi-commodity case. Proving convergence for any choice of the scalar functions $\lambda_i(\cdot)$ using the conditional entropy as a Lyapunov function we would need the property $f_{|\mathcal{P}_i} \cdot \ell(\tilde{f})_{|\mathcal{P}_i} \leq \tilde{f}_{|\mathcal{P}_i} \cdot \ell(\tilde{f})_{|\mathcal{P}_i}$ for every commodity $i \in [k]$ independently. This, however, does not hold. In order to prove convergence in the multi-commodity case we will use the Beckmann-McGuire-Winsten potential [11] as defined in Equation (1.8):

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} \ell_e(u)\, du \quad .$$

We have seen that Wardrop equilibria minimise this potential function [11]. We start by computing the derivative with respect to time of Φ.

Lemma 2.3. *For the replicator dynamics (2.2), the time derivative of the potential is given by*

$$\dot{\Phi}(f) = \sum_{i \in [k]} \lambda_i(f) \left(L_i(f)^2 - \sum_{P \in P_i} \frac{f_P}{r_i} (\ell_P(f_e))^2 \right) \leq 0 \quad .$$

Proof. The derivative of Φ with respect to time is

$$\dot{\Phi}(f) = \sum_{e \in E} \dot{f}_e \cdot \ell_e(f_e) = \sum_{e \in E} \sum_{P \ni e} \dot{f}_P \cdot \ell_e(f_e) \quad .$$

Substituting the replicator dynamics (2.2) into this we obtain

$$\begin{aligned}
\dot{\Phi}(f) &= \sum_{e \in E} \sum_{P \ni e} (\lambda_{i_P}(f) \cdot f_P \cdot (L_{i_P}(f) - \ell_P(f))) \cdot \ell_e(f_e) \\
&= \sum_{i \in [k]} \lambda_i(f) \sum_{P \in P_i} \sum_{e \in P} f_P \cdot (L_i(f) - \ell_P(f)) \cdot \ell_e(f_e) \\
&= \sum_{i \in [k]} \lambda_i(f) \sum_{P \in P_i} f_P \cdot (L_i(f) - \ell_P(f)) \cdot \ell_P(f) \\
&= \sum_{i \in [k]} \lambda_i(f) \left(L_i(f) \sum_{P \in P_i} f_P \cdot \ell_P(f) - \sum_{P \in P_i} f_P \cdot \ell_P(f)^2 \right) \\
&= \sum_{i \in [k]} \lambda_i(f) \left(r_i \cdot L_i(f)^2 - \sum_{P \in P_i} f_P \cdot (\ell_P(f))^2 \right) \\
&= \sum_{i \in [k]} \lambda_i(f) \cdot r_i \cdot \left(L_i(f)^2 - \sum_{P \in P_i} \frac{f_P}{r_i} (\ell_P(f))^2 \right) \quad .
\end{aligned}$$

By Jensen's inequality, each term of this sum is non-positive. □

Convergence follows directly from this lemma.

Theorem 2.4. *Consider a multi-commodity instance* Γ*. For* $i \in [k]$*, let* $\lambda_i(\cdot)$ *be a continuous, positive function. Let* $\xi(f_0, t)$ *be a solution to the replicator dynamics (2.2) for some initial population* $f_0 \in \mathrm{int}(\mathcal{F}_\Gamma)$*. Then,* $\xi(f_0, t)_{t\to\infty} \to \mathcal{F}_\Gamma^{NE}$*.*

Proof. We use the potential Φ as a Lyapunov function. By Jensen's inequality, every term of the sum

$$\dot{\Phi}(\xi) = \sum_{i\in[k]} \lambda_i(\xi) \cdot r_i \cdot \left(L_i(\xi)^2 - \sum_{P\in P_i} \frac{\xi_P}{r_i} (\ell_P(\xi))^2 \right)$$

is non-positive and zero if ξ is at a Wardrop equilibrium. Furthermore, for any $\xi \in \mathrm{int}(\Delta)$, $\dot{\Phi}(\xi)$ can be zero only if for all $i \in [k]$ and all paths $P \in \mathcal{P}_i$ we have $\ell_P(\xi) = L_i(\xi)$. Hence, $\Phi(f) - \Phi^*$ is a Lyapunov function satisfying the conditions of Theorem 1.7, and furthermore, in $\mathrm{int}(\Delta)$ the vector field of the replicator dynamics (2.1) points strictly into every contour set of Φ and hence $\xi(f_0, t)_{t\to\infty} \to \mathcal{F}_\Gamma^{NE}$. □

2.3 Approximate Equilibria and Convergence Time

In the previous section we have established convergence towards Wardrop equilibria in the long run. An interesting question is now: How long does it take for the population to reach an equilibrium? We believe that this is an issue of particular importance since equilibria are only meaningful if they can be reached within reasonable time.

Since the population cannot reach an *exact* equilibrium (unless starting at an equilibrium), we consider approximate equilibria. We can define approximate equilibria as population states in which the latencies sustained by the agents are close to the respective average latencies L_i rather than equal to L_i. Note, however, that we cannot require this property for *all* agents since the population on a path with constant large latency will never vanish completely. Furthermore, we cannot expect the replicator dynamics to come close to the global Wardrop equilibrium in terms of Euclidean distance or in terms of the potential Φ since we can make the time to come close to the global Wardrop equilibrium arbitrarily large by making the initial flow on a path that is used in the Wardrop equilibrium arbitrarily small. The idea behind the following definition of approximate equilibria is not to wait for such minorities.

Definition 2.2 (ϵ-equilibrium)**.** *For a single-commodity instance* Γ *and a flow* $f \in \mathcal{F}_\Gamma$*, let* $\mathcal{P}_\epsilon$ *denote the set of paths with latency above* $(1+\epsilon)\cdot L(f)$*, i. e.,* $\mathcal{P}_\epsilon = \{P \in \mathcal{P} \mid \ell_P(f) > (1+\epsilon)L(f)\}$*. Furthermore, let* $f_\epsilon = \sum_{P\in\mathcal{P}_\epsilon} f_P$

*be the total volume of agents utilising such paths. Then, the flow f is at an ϵ-*equilibrium *if and only if $f_\epsilon \leq \epsilon \cdot r$.*

Approximate equilibria in this sense can be interpreted informally as populations where *almost all agents are almost satisfied*. In the following, we give bounds on the time until the replicator dynamics reaches an ϵ-equilibrium for the first time. Let us point out, however, that the dynamics does not necessarily stay at an ϵ-equilibrium once an ϵ-equilibrium is reached. A population may leave an ϵ-equilibrium after an arbitrarily long period of time. This may happen, e. g., if an initially small population share utilising a high-latency path starts to grow to above ϵ. The opposite may also happen. Population shares on paths with low latency may grow, thus lowering the average and making agents on other paths become dissatisfied again. In this sense, ϵ-equilibria are *transient*.

Recall that in order for time bounds to be meaningful, we must now make a choice for the scalar functions $\lambda_i(\cdot)$ in Equation (2.2) which determine the speed at which the solution orbit is traversed. This choice should make the speed independent of the scale by which we measure latency. In Section 1.2.1 we have seen that $\lambda_i(f) = 1/L_i(f)$ arises naturally from the motivation by *fitness-proportional replication*. Therefore, throughout this section, we will fix $\lambda_i(f) = 1/L_i(f)$.

We can now prove bounds on the convergence time. For single-commodity instances, our bounds indicate that the replicator dynamics converges towards approximate equilibria quickly. This bound is essentially tight as is shown by a corresponding lower bound. For multi-commodity instances the upper bound is slightly weaker.

2.3.1 Upper Bound for Single-Commodity Instances

Our bounds is in terms of the ratio between maximal and optimal latency. Recall that $L_{\max} = \max_{f \in \mathcal{F}_\Gamma} L(f)$ and $L_{\min} = \min_{f \in \mathcal{F}_\Gamma} L(f)$ denote the maximum and minimum total latency, respectively.

Theorem 2.5. *For any solution of the replicator dynamics (2.1) for a single-commodity instance Γ and any initial population $f_0 \in \mathcal{F}_\Gamma$, the total amount of time that the population is not at an ϵ-equilibrium is bounded by*

$$\mathcal{O}\left(\frac{1}{\epsilon^3} \cdot \ln\left(\frac{L_{\max}}{L_{\min}}\right)\right) .$$

Proof. We derive a lower bound on the rate by which the potential Φ decreases as long as we are not at an ϵ-equilibrium. In order to obtain a positive lower bound on the potential, we consider the shifted potential

$$\Phi'(f) = \Phi(f) + L_{\min} \cdot r$$

instead. Since $L_{\min} \cdot r$ is constant, we have $\dot{\Phi}'(f) = \dot{\Phi}(f)$. As we know from Lemma 2.3,

$$\dot{\Phi}'(f) = \dot{\Phi}(f) = \lambda(f) \cdot r \cdot \left(L(f)^2 - \sum_{P \in \mathcal{P}} \frac{f_P}{r} (\ell_P(f))^2 \right) \quad . \tag{2.5}$$

When is the absolute value of this term minimal given that at least a volume of $\epsilon \cdot r$ agents utilises paths with latency at least $(1+\epsilon) \cdot L(f)$? As a consequence of Jensen's inequality, for any fixed $L(f)$, the term $\sum_{P \in \mathcal{P}} (f_P/r) \cdot \ell_P(f)^2$ is minimised when the remaining volume of $r \cdot (1-\epsilon)$ agents utilise paths with equal latency which we denote by $\tilde{\ell}$. Then, $L(f) \geq \epsilon \cdot (1+\epsilon) \cdot L(f) + (1-\epsilon) \cdot \tilde{\ell}$ and hence

$$\tilde{\ell} \leq \frac{1-\epsilon-\epsilon^2}{1-\epsilon} \cdot L(f) = \left(1 - \frac{\epsilon^2}{1-\epsilon}\right) \cdot L(f) \tag{2.6}$$

Combining this with Equation (2.5) we obtain

$$\dot{\Phi}'(f) \leq \lambda(f) \cdot r \cdot \left(L(f)^2 - \left(\epsilon \cdot ((1+\epsilon) \cdot L(f))^2 + (1-\epsilon) \cdot \tilde{\ell}^2\right)\right) \quad .$$

Substituting $\tilde{\ell}$ according to Equation (2.6) we get

$$\begin{aligned} \dot{\Phi}'(f) &\leq -\lambda(f) \cdot r \cdot L(f)^2 \cdot \frac{\epsilon^3}{1-\epsilon} \\ &\leq -\epsilon^3 \cdot r \cdot L(f) \ , \end{aligned} \tag{2.7}$$

where we have used $\lambda(f) = 1/L(f)$. Now observe that

$$\begin{aligned} L(f) &= \sum_{e \in E} \frac{f_e}{r} \ell_e(f_e) \\ &\geq \frac{1}{r} \cdot \sum_{e \in E} \int_0^{f_e} \ell_e(u)\, du \\ &= \frac{\Phi(f)}{r} \end{aligned}$$

by monotonicity of the ℓ_e. Hence, by definition of $L_{\min}$ we have $2\,L(f) \geq L_{\min} + \Phi(f)/r = \Phi'(f)/r$. Finally, substituting this into Equation (2.7) we obtain the differential inequality

$$\dot{\Phi}'(f) \leq -\frac{\epsilon^3}{2}\, \Phi'(f)$$

which implies that any solution that is not at an ϵ-equilibrium during the interval $(t, t+\tau)$ satisfies

$$\Phi'(f(t+\tau)) \leq \Phi'(f(t)) \cdot \exp\left(-\frac{\epsilon^3}{2} \cdot \tau\right) \quad .$$

For $i \geq 1$ let $(t_i, t_i + \tau_i)$ denote the intervals at which the population is not at an ϵ-equilibrium. (These intervals are open due to the requirement that $f_\epsilon > r \cdot \epsilon$. The sole exception may be the first interval $[0, \tau_1)$ if $t_1 = 0$.) By induction on i,

$$\Phi'(f(t_i + \tau_i)) \leq \Phi'(f(t_1)) \cdot \exp\left(-\frac{\epsilon^3}{2} \cdot \sum_{j=1}^{i} \tau_j\right) \quad .$$

Let $T = \sum_{j=1}^{\infty} \tau_i$ denote the total amount of time that the population is not at an ϵ-equilibrium. Since Φ' is, by definition, bounded from below by $L_{\min} \cdot r$, this implies that

$$\Phi'(f_0) \cdot \exp\left(-\frac{\epsilon^3}{2} \cdot T\right) \geq L_{\min} \cdot r \;\; ,$$

or, equivalently

$$T \leq \frac{2}{\epsilon^3} \cdot \ln\left(\frac{\Phi'(f_0)}{L_{\min} \cdot r}\right) \quad .$$

Since $\Phi'(f_0) \leq 2 \cdot r \cdot L_{\max}$, this establishes our assertion. $\square$

2.3.2 Upper Bound for Multi-Commodity Instances

In the asymmetric case our bounds on the convergence time are weaker. Our definition of approximate equilibria can be generalised to multi-commodity instances in the following way.

Definition 2.3 (ϵ-equilibrium)**.** *For a multi-commodity instance Γ and a population $f \in \mathcal{F}_\Gamma$, let $\mathcal{P}_\epsilon = \bigcup_{i \in [k]} \{P \in \mathcal{P}_i \mid \ell_P(f) > (1+\epsilon) L_i(f)\}$. Furthermore, let $f_\epsilon = \sum_{P \in \mathcal{P}_\epsilon} f_P$ be the total volume of agents utilising such paths. Then, a flow f is at an* ϵ-equilibrium *if and only if $f_\epsilon \leq r \cdot \epsilon$.*

Our bound is in terms of the ratio between maximum average latency and minimum average latency of a commodity. Recall that we have defined $\tilde{L}_{\min} = \min_{i \in [k]} \min_{f \in \mathcal{F}_\Gamma} L_i(f)$.

Theorem 2.6. *The replicator dynamics (2.2) for a multi-commodity instance Γ and any initial population $f_0 \in \mathcal{F}_\Gamma$ the total amount of time that the population is not at an ϵ-equilibrium is bounded by*

$$\mathcal{O}\left(\frac{1}{\epsilon^3} \cdot \frac{L_{\max}}{\tilde{L}_{\min}}\right) \quad .$$

Proof. The proof goes along the lines of the proof of Theorem 2.5. In the multi-commodity case, by Lemma 2.3,

$$\dot{\Phi}(f) = \sum_{i \in [k]} \lambda_i(f) \cdot r_i \cdot \left(L_i(f) - \sum_{P \in \mathcal{P}_i} \frac{f_P}{r_i} \ell_P(f)^2\right) \quad .$$

Let ϵ_i denote the volume of agents in commodity i that utilise paths with latency above $(1+\epsilon)L_i$ and note that $\sum_{i\in[k]} \epsilon_i > \epsilon \cdot r$ as long as we are not at an ϵ-equilibrium. With the same argument that led to Equation (2.6), the contribution of commodity i to the above expression for the potential gain is minimal if the remaining agents utilise paths with equal latency $\tilde{\ell}_i$. Since $L_i(f) \geq (r_i - \epsilon_i)/r_i \cdot (1+\epsilon)\, L_i + \epsilon_i/r_i \cdot \tilde{\ell}_i$,

$$\tilde{\ell}_i \leq L_i(f) \cdot \left(1 - \frac{\epsilon_i \cdot \epsilon}{r - \epsilon_i}\right) \quad .$$

Substituting this into our expression for $\dot{\Phi}(f)$, the contribution of commodity i to the potential gain is then

$$\dot{\Phi}_i(f) = -\lambda_i(f) \cdot r_i \cdot L_i(f)^2 \cdot \frac{\epsilon^2 \cdot \epsilon_i}{r_i - \epsilon_i} \leq -L_i(f) \cdot \epsilon^2 \cdot \epsilon_i \quad .$$

Summing over all commodities and using the pessimistic estimate $L_i(f) \geq \tilde{L}_{\min}$, we obtain

$$\dot{\Phi}(f) \leq -\sum_{i\in[k]} L_i(f) \cdot \epsilon^2 \cdot \epsilon_i \leq -\tilde{L}_{\min} \cdot \epsilon^3 \cdot r$$

implying that for any interval $(t, t+\tau)$ during which the population is not at a ϵ-equilibrium,

$$\Phi(f(t+\tau)) \leq \Phi(f(t)) - \tau \cdot r \cdot \epsilon^3 \cdot \tilde{L}_{\min} \quad .$$

Let T denote the total amount of time at which the population is not at an ϵ-equilibrium. Since Φ is bounded from below by 0, it holds that $\Phi(f_0) - T\, r\, \epsilon^3\, \tilde{L}_{\min} \geq 0$ or, equivalently,

$$T \leq \frac{\Phi(f_0)}{\epsilon^3 \cdot r \cdot \tilde{L}_{\min}} \leq \frac{L_{\max}}{\epsilon^3 \cdot \tilde{L}_{\min}}$$

which is our claim. □

2.3.3 Lower Bound for Single-Commodity Instances

In this section we give a lower bound on the time to reach an ϵ-equilibrium for the single-commodity case. Our bound is tight with respect to Theorem 2.6 up to a factor of $\mathcal{O}\left(\epsilon \cdot \ln(\epsilon^{-1})\right)$.

Theorem 2.7. *For every $r, \epsilon > 0$ there exists a single-commodity singleton instance Γ with $\ell_{\max}/\ell_{\min} = r$ and an initial flow vector $f_0 \in \mathcal{F}_\Gamma$ such that for the solution $\xi(f_0, t)$ of the replicator dynamics (2.1), $f(t)$ is not at an ϵ-equilibrium for any $t \in [0, \tau]$ with*

$$\tau = \Omega\left(\frac{1}{\epsilon^2} \cdot \ln\left(\frac{1}{\epsilon}\right) \cdot \ln\left(\frac{\ell_{\max}}{\ell_{\min}}\right)\right) \quad .$$

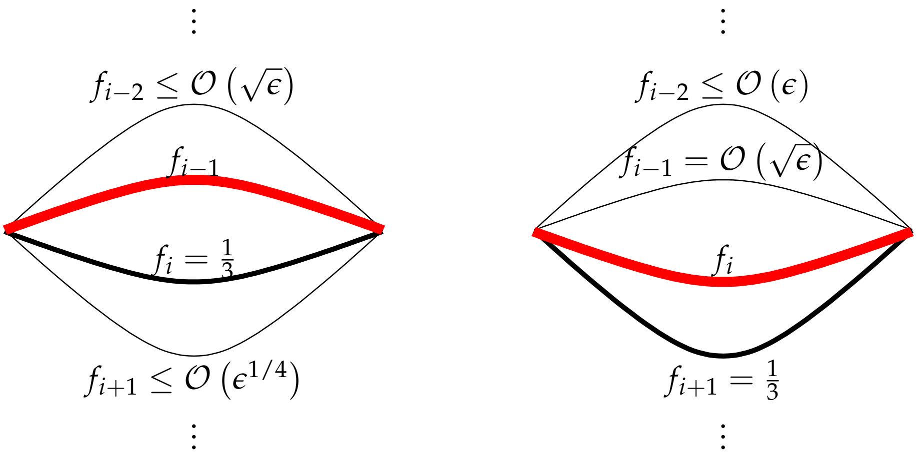

Figure 2.1: Flow conditions at the beginning (left) and at the end (right) of phase i. Agents on link $i-1$ are unsatisfied.

Proof. We construct a network of $n+1$ parallel links $0, \ldots, n$ with constant latency functions

$$\ell_i(x) = \frac{1}{(1+4\,\epsilon)^i}$$

Since $\ell_i(\cdot)$ is constant we simply write ℓ_i rather than $\ell_i(x)$. We consider n phases $1, \ldots, n$ of length

$$T = \frac{\ln(1/(3\,\epsilon))}{8\,\epsilon} \quad ,$$

and $t_i = (i-1) \cdot T$ denotes the start of phase i. We construct an initial flow vector f_0 such that the flow cascades down from link 0 to n where in phase i, the agents on link $i-1$ are unsatisfied (in that $\ell_{i-1} > (1+\epsilon)L$) and $f_{i-1} \geq \epsilon$. Throughout the phase the majority of the agents utilises links $i-1$, i, and $i+1$ whereas the flow on other links is negligible. We have

$$r = \frac{\ell_{\max}}{\ell_{\min}} = (1+4\,\epsilon)^n$$

and hence for any given value of r, our construction has

$$n = \log_{(1+4\epsilon)}(r) \geq \frac{1}{4\,\epsilon} \ln(r)$$

phases. Since the length of each phase is $\Theta(\ln(\epsilon^{-1})/\epsilon)$, the total time to reach an ϵ-equilibrium is $\Omega(\epsilon^{-2}\ln(\epsilon^{-1})\ln(r))$, our desired bound.

We start with two observations:

(O1) For constant latency functions, $\dot{L}(f) < 0$ for all $f \in \mathcal{F}_\Gamma$.

(O2) As long as $L(f) > \ell_i$, we have $\dot{f}_i > 0$, and as soon as $L(f) < \ell_i$, we have $\dot{f}_i < 0$.

We construct our initial population such that for phase $i \in \{1, \ldots, n\}$ the following properties hold (see Figure 2.1):

(P1) For every link $j \leq i-2$, $f_j(t_i) \leq \epsilon^{(i-j-1)/2}$.

(P2) For every link $j \geq i+1$, $f_j(t_i) \leq (3\,\epsilon)^{(j-i)/4}$.

(P3) $f_i(t_i) \geq \frac{1}{3}$.

(P4) $f_i(t_i) \leq \frac{1}{3}$.

Claim 2.8. *Consider a phase $i \in [n]$ and a flow vector $\hat{f} \in \mathcal{F}_\Gamma$ that satisfies $\hat{f}_j \leq (3\,\epsilon)^{(j-i)/4}$ for $j \geq i+2$. Then, there exists an initial flow vector f_0 such that properties (P1) – (P4) are satisfied in phases $1, \ldots, i$ and $f_j(t_i) = \hat{f}_j$ for any $j \geq i+1$. Furthermore, in any phase $j \in \{1, \ldots, i\}$, the agents on link $i-1$ are unsatisfied, and $f_{i-1}(t) \geq \epsilon$ for any $t \in [t_i, t_{i+1})$. Thus, the population is never at an ϵ-equilibrium throughout phases $1, \ldots, i$.*

Note that the requirements for $\hat{f}$ are compatible with (P2) in phase $i+1$. For $i = n$, this claim implies the assertion of the theorem.

Proof. The proof is by induction on i. The properties can clearly be satisfied for the first phase $i = 1$. We will see below that satisfying the four properties implies that the population is not at an ϵ-equilibrium throughout phase i.

For the induction step, fix a flow vector $\hat{f}$ that satisfies the conditions of Claim 2.8 for round $i+1$. First note that in round $i+1$ properties (P2), (P4) as well as $f_j(t_{i+1}) \leq \hat{f}_j$ for $j \geq i+2$ can be satisfied by choosing $f_j(t_i)$ sufficiently small for $j \geq i+1$. We will see that they can also be chosen large enough respecting property (P2) in round i such that $f_j(t_{i+1}) \geq \hat{f}_j$ for $j \geq i+2$ and $f_{i+1}(t_{i+1}) \geq 1/3$.

We start with upper and lower bounds on L to deduce bounds on the growth rates of the population shares. Consider a flow $f(t_i)$ at the beginning of phase i and assume that all properties are satisfied.

Let $f(t_i)$ denote the flow at the beginning of phase i and let $\tilde{f}$ denote the flow obtained from $f(t_i)$ by shifting agents from links $i+1, \ldots, n$ to link $i-1$. Then, $L(f(t_i)) \leq L(\tilde{f})$. By (P3) there is a flow volume of $1/3$ on link i and hence, $\tilde{f}_{i-1} \leq 2/3 - \sum_{j=0}^{i-2} f_j(t_i)$.

$$\begin{aligned}
L(f(t_i)) &\leq L(\tilde{f}) \\
&\leq \frac{1}{3}\ell_i + \left(\frac{2}{3} - \sum_{j=0}^{i-2} f_j(t_i)\right) \ell_{i-1} + \sum_{j=0}^{i-2} f_j(t_i)\,\ell_j \\
&= \ell_{i-1}\left(\frac{1}{3}\cdot\frac{1}{1+4\epsilon} + \frac{2}{3} - \sum_{j=0}^{i-2} f_j(t_i) + \sum_{j=0}^{i-2} f_j(t_i)\frac{\ell_j}{\ell_{i-1}}\right) .
\end{aligned}$$

Using the bounds on f_j for $j \leq i-2$ specified by (P1) we obtain

$$L(f(t_i)) \leq \ell_{i-1}\left(\frac{1}{3}\cdot\frac{1}{1+4\epsilon} + \frac{2}{3} - \sum_{j=1}^{\infty}\sqrt{\epsilon}^j + \sum_{j=1}^{\infty}\left(\sqrt{\epsilon}(1+4\epsilon)\right)^j\right)$$

Since $\sum_{i=1}^{\infty} X^j = X/(1-X)$ for $X < 1$, we have

$$\begin{aligned}
L(f(t_i)) &= \ell_{i-1}\left(1 - \frac{4\epsilon}{3(1+4\epsilon)} - \frac{\sqrt{\epsilon}}{1-\sqrt{\epsilon}} + \frac{\sqrt{\epsilon}(1+4\epsilon)}{1-\sqrt{\epsilon}(1+4\epsilon)}\right) \\
&= \ell_{i-1}\left(1 - \frac{4\epsilon}{3(1+4\epsilon)} + \mathcal{O}\left(\epsilon^{3/2}\right)\right) \\
&\leq \ell_{i-1}(1-\epsilon) \qquad (2.8)
\end{aligned}$$

if ϵ is small. This holds at the beginning of phase i and by observation (O1) also throughout the entire phase.

For the lower bound we consider the value of $L(f(t_{i+1}))$ which by observation (O1) is also a lower bound on L throughout phase i. Note that the amount of flow we have to assign to the links $i+2,\ldots,n$ due to the requirement that $f(t_{i+1})$ matches $\hat{f}$ on these links still respects property (P2) in phase $i+1$. Fix such a flow vector and consider the modified flow vector $\tilde{f}$ obtained by shifting all agents on links $0,\ldots,i-1$ to link i. Then, $L(\tilde{f}) \leq L(f(t))$ for $t \in [t_i, t_{i+1}]$. Consequently,

$$\begin{aligned}
L(f(t)) &\geq L(\tilde{f}) \\
&\geq \frac{1}{3}\ell_{i+1} + \left(\frac{2}{3} - \sum_{j=i+2}^{n}(3\,\epsilon)^{(j-(i+1))/4}\right)\ell_i + \sum_{j=i+2}^{n}(3\,\epsilon)^{(j-(i+1))/4}\ell_j \\
&\geq \frac{1}{3}\ell_{i+1} + \left(\frac{2}{3} - \sum_{j=1}^{\infty}(3\,\epsilon)^{j/4}\right)\ell_i + \ell_{i+1}\sum_{j=1}^{\infty}\left(\frac{(3\,\epsilon)^{j/4}}{1+4\epsilon}\right)^j \\
&= \ell_{i+1}\left(1 + \frac{8}{3}\epsilon - \frac{(1+4\epsilon)(3\,\epsilon)^{1/4}}{1-(3\,\epsilon)^{1/4}} + \frac{(3\,\epsilon)^{1/4}/(1+4\epsilon)}{1-(3\,\epsilon)^{1/4}/(1+4\epsilon)}\right) \\
&= \ell_{i+1}\left(1 + \frac{8}{3}\epsilon - \mathcal{O}\left(\epsilon^{5/4}\right)\right) \\
&\geq \ell_{i+1}(1+2\,\epsilon) \qquad (2.9)
\end{aligned}$$

if ϵ is small. This holds at the end of phase i and consequently by observation (O1) throughout the entire phase i.

Equation (2.8) implies that $\ell_{i-1} > L \cdot (1+\epsilon)$ and hence the agents on link $i-1$ are unsatisfied in the entire phase i. We show that there are at least ϵ such agents on link $i-1$ throughout the phase by using Equation (2.9):

$$\begin{aligned}
\ell_{i-1} &= (1+4\epsilon)^2 \ell_{i+1} \\
&\leq L \cdot \frac{(1+4\epsilon)^2}{1+2\,\epsilon} \\
&= L \cdot \left(1 + \frac{6\epsilon + 16\epsilon^2}{1+2\,\epsilon}\right) \\
&\leq L \cdot (1+8\epsilon)
\end{aligned}$$

if ϵ is small. Substituting this into the replicator dynamics (2.1), we see that $\dot{f}_{i-1} \geq -8\,\epsilon\, f_{i-1}$. Due to the upper bounds imposed by our properties on $f_j(t_i)$ for $j \neq i-1$ and since the total flow is 1, the remaining flow volume is on link $i-1$, i. e., $f_{i-1}(t_i) \geq 1/3$. Then,

$$f_{i-1}(t_i + t) \geq \frac{1}{3} \cdot \exp(-8\,\epsilon\, t)$$

which in turn implies $f_{i-1}(t_i + t) \geq \epsilon$ as long as

$$t \leq \frac{\ln\left(\frac{1}{3\,\epsilon}\right)}{8\,\epsilon} = T \ .$$

Consequently, a volume of at least ϵ agents utilises link $i-1$ in phase i, and these agents are unsatisfied. Hence, the population is not at an ϵ-equilibrium throughout phase i.

We now show that property (P1) is satisfied at the beginning of the next phase $i+1$. Consider the upper links $0, \ldots, i-2$. From Equation (2.8) it follows that

$$\begin{aligned}
L(f(t_i)) &\leq \ell_{i-1}\,(1-\epsilon) \\
&= \ell_{i-2} \cdot \frac{1-\epsilon}{1+4\,\epsilon} \\
&\leq \ell_{i-2} \cdot \left(1 - \frac{9}{2}\,\epsilon\right) \ .
\end{aligned}$$

if ϵ is small. Hence, for all links $j \leq i-2$ we have $\dot{f}_j \leq -\frac{9}{2}\,\epsilon\, f_j$ implying that at the end of the phase,

$$\begin{aligned} f_j(t_{i+1}) = f_j(t_i + T) &\leq f_j(t_i) \cdot \exp\left(-\frac{9}{2}\,\epsilon\, T\right) \\ &= f_j(t_i) \cdot \exp\left(-\frac{9}{16}\ln(1/(3\,\epsilon))\right) \\ &= f_j(t_i) \cdot (3\,\epsilon)^{\frac{9}{16}} \\ &\leq f_j(t_i) \cdot \sqrt{\epsilon} \end{aligned}$$

if ϵ is small. This implies that these links satisfy property (P1) for the phase $i+1$.

For the bottom links $j \geq i+1$, Equation (2.9) implies that we have a growth rate of $\dot{f}_j \geq 2\,\epsilon\, f_j$ in phase i. Hence, for $j \geq i+1$,

$$\begin{aligned} f_j(t_{i+1}) &\geq f_j(t_i) \cdot \exp(2\,\epsilon\, T) \\ &= f_j(t_i) \cdot \exp\left(2\,\epsilon \cdot \frac{\ln\left(\frac{1}{3\epsilon}\right)}{8\,\epsilon}\right) \\ &= f_j(t_i) \cdot \left(\frac{1}{3\,\epsilon}\right)^{\frac{1}{4}} . \end{aligned} \tag{2.10}$$

We can now show that we can satisfy property (P3) in phase $i+1$ and ensure that $f_j(t_{i+1})$ matches $\hat{f}$ on links $j \geq i+2$. First note that due to Equation (2.10), $f_{i+1}(t_i) \geq (3\,\epsilon)^{1/4}$ implies $f_{i+1}(t_{i+1}) \geq 1/3$. Second, consider a flow vector $\hat{f}$ that satisfies the condition of Claim 2.8, $\hat{f}_j \leq (3\,\epsilon)^{(j-i-1)/4}$ for $j \geq i+2$. Then, due to Equation (2.10), $f_j(t_i) \geq (3\,\epsilon)^{(j-i)/4}$ implies that $f_j(t_{i+1}) \geq \hat{f}_j$.

Altogether, since for $j \geq i+1$, $f_j(t_{i+1})$ is continuous and strictly increasing in $f_j(t_i)$ we can satisfy the above inequalities with equality by choosing the values of the $f_j(t_i)$ appropriately. More precisely, there exists a flow vector $\hat{f}'$ satisfying $f'_j \leq (3\,\epsilon)^{(j-i)/4}$ for $j \geq i+1$ such that $f_j(t_i) = \hat{f}'_j$ for $j \geq i+1$ implies that $f_{i+1}(t_{i+1}) = 1/3$ and $f_j(t_{i+1}) = \hat{f}_j$ for $j \geq i+2$. Since $\hat{f}'$ is compatible with the requirements of the induction hypothesis, we can ensure that $f_j(t_i) = \hat{f}'_j$ for $j \geq i+1$. □

The proof of the claim concludes the proof of the theorem. □

2.4 Summary

We have seen that Wardrop equilibria satisfy a stronger stability criterion than being at a Nash equilibrium alone: They are evolutionarily stable. Whereas this is an appealing property in itself, it also enables us to show that Wardrop equilibria are asymptotically stable in the replicator dynamics. Moreover, the set of Wardrop equilibria attracts the entire interior of the set of feasible flows. In addition, we have given bounds on the time to reach approximate Wardrop equilibria.

Our results show that a population of selfish agents is able to attain an equilibrium state in the long run by following simple rules requiring minimal rationality and knowledge assumptions only. This is particularly important for scenarios that are typically modelled as selfish routing games, since these scenarios often fail to satisfy stronger assumptions. These results strengthen the concept of Wardrop equilibria as solutions of selfish routing games.

The single assumption that could not be abandoned to obtain the stability and convergence properties presented in this chapter is the assumption of perfect information. Our analyses require that all rerouting decisions made by other agents become visible immediately. In the next chapter we consider a model in which this is not the case and show how convergence can be guaranteed nevertheless.

Chapter 3

Coping with Stale Information

In the previous chapter we have seen that, in the long run, a population of selfish agents is able to attain a Wardrop equilibrium in a distributed way. However, our dynamic population model treats the aspect of concurrency in an idealistic fashion. It is implicitly assumed that any effect that the migration of population shares has on latency can be observed without a delay. In this sense, we have still assumed perfect information. A reasonable model of concurrency, however, should take into account the possibility that information may be stale: In practise, e. g., a rerouted flow may unfold its full volume only after a re-initialisation phase, and "observing" the latency of a path may actually comprise measurements taken over a certain period of time. In such a scenario, naive policies like the best response dynamics may lead to overshooting and oscillation effects. These effects are well-known in practise and may seriously harm performance [47, 49, 62, 73]. In fact, oscillation effects are a major obstacle in designing real-time dynamic traffic engineering protocols. Similar negative results have been obtained in theory [17, 52].

The strength of these oscillation effects chiefly depends on three quantities: the maximum age of the information, the sensitivity of the latency functions to small changes of the flow, and the speed at which the rerouting policy shifts flow. Whereas the first two parameters are determined by the system, the third is controlled by the rerouting policy, e. g., in case of the replicator dynamics, by the scalar function $\lambda_i(f)$ in Equation (2.2).

In this chapter, we ask the question of how agents should behave in order to guarantee convergence to Wardrop equilibria despite the fact that information may be stale. As a model of stale information we consider the bulletin board model introduced by Mitzenmacher [52] in the context of dynamic load balancing. In this model, all information relevant to the routing process is posted on a bulletin board and updated at regular intervals of length T. The information on the bulletin board is accessible to all agents. Again, agents may reroute their traffic at points in time that are the arrival dates of a Poisson process. We will see that in the bulletin board model, naive rerouting policies like the best response policy oscillate and fail to converge.

We consider a class of simple, stateless rerouting policies which also contains the replicator dynamics discussed in the preceding chapter and devise sufficient properties for dynamics from this class to converge. More precisely, we show that for any class of latency functions with bounded slope and any finite update period length T there exist policies that converge towards the set of Wardrop equilibria. We also give bounds on the time to converge towards approximate equilibria.

3.1 Related Work

The bulletin board model of stale information was introduced by Mitzenmacher [52] in a different context. In Mitzenmacher's model, jobs (agents) enter the system at Poisson rates, are assigned to machines (resources) and are removed from the system after being processed. Both the number of agents and the number of resources go to infinity whereas the ratio between the number of jobs and the number of machines is fixed. In our model the number of resources (the edges of the considered network) is fixed whereas the number of agents is infinite. Agents are not removed from the system but merely reassign their load at intervals. Mitzenmacher's results are mainly negative: Decisions based on stale information can degrade performance even in comparison with a fully random assignment of jobs. For some naive strategies this is the case in our model as well. However, as a positive result we can show that the system can come to an equilibrium solely based on old information.

Knowing that information may be stale, it is natural to take the age of the information into account for the decision process. In fact, in [22] it is shown that this increased computational effort can improve network performance. For the task of job scheduling, a natural approach is to assign a job to the machine with expected shortest queue. Such a strategy is analysed by Cao and Nyberg [17]. They show that this does not always minimise the average waiting time. In our scenario, the age of the information is not known to the agents.

In classical optimisation one considers the minimisation of a global cost function rather than the computation of a Wardrop equilibrium. As we have seen, the two problems can be cast into one another by using marginal cost latency functions [70, 83]. For example, Bertsekas and Tsitsiklis [13] consider a distributed algorithm for optimising a global cost function with asynchronous agents. In game theoretic terminology, their algorithm is comparable to a careful best response policy. The computational effort required by this algorithm is greater than the negligible computational effort of our sampling policies. In the model of Bertsekas and Tsitsiklis, traffic is rerouted at discrete time steps whereas our model is continuous. Furthermore, we also give bounds on the time of convergence based on network parameters.

In evolutionary game theory, similar problems are faced when considering discretisations of the replicator dynamics. E. g., the Dekel-Scotchmer Example [23] shows that in discrete versions of the replicator dynamics, it is not even guaranteed that dominated strategies vanish.

The problem of oscillation due to stale information is also well-known in practise. We refer to the respective section of Chapter 5 for a discussion of the literature.

Our model of stale information will lead to a system of differential equations which depend on retarded time. Generalisations of Lyapunov's method for dynamical systems of a similar kind can be found in [24].

3.2 Smooth Rerouting Policies and Stale Information

The traffic model considered in this chapter is again the Wardrop model presented in Section 1.3. As has been mentioned, our main result depends on the maximum slope of the latency functions, which in turn depends on the scale by which we measure latency and the scale by which we measure flow. For the sake of simplicity, throughout this chapter, we normalise the total flow demand to be $r = 1$. Under this normalisation, we can express our results in terms of the maximum slope and the maximum path latency $\ell_{\max} = \max_{P\in\mathcal{P}} \sum_{e\in P} \ell_e(1)$

3.2.1 Smooth Rerouting Policies

In this chapter we consider a very broad class of adaptive routing policies. These policies can be described in the following way. Each agent becomes active at discrete points in time generated at fixed Poisson rates for each agent independently. Whenever an agent is activated, she revises her routing strategy (i. e. the selected path) by performing the following two steps. Consider an agent of commodity $i \in [k]$ currently using path $P \in \mathcal{P}_i$.

1. *Sampling:* The agent samples a path $Q \in \mathcal{P}_i$ according to some probability distribution. Let $\sigma_{PQ}(f)$ denote the probability that an agent currently using path P samples path Q.

2. *Migration:* The agent switches to the sampled path Q with probability $\mu(\ell_P, \ell_Q)$.

We illustrate the class of policies in this framework by giving some examples for possible sampling and migration rules.

Typically, $\sigma_{PQ}(f)$ will be independent of the originating path P. The simplest sampling rule is the *uniform sampling* rule which assigns uniform probabilities to all paths, i. e., $\sigma_{PQ}(f) = 1/|\mathcal{P}_i|$ for all $i \in [k]$ and $P, Q \in \mathcal{P}_i$. Instead of choosing paths uniformly at random, one can also imagine that agents sample other agents from the same commodity uniformly at random as we have done for the replicator dynamics. We call this approach *proportional sampling* since the probability to sample path Q is proportional to the fraction of agents using it, i. e. $\sigma_{PQ}(f) = f_Q/r_i$.

In order to guarantee that unique solutions of the dynamic system described by the above rules exist, we assume that $\sigma_{PQ}(f)$ is a Lipschitz-continuous function of f. Also note that a population cannot converge towards a Wardrop equilibrium if $\sigma_{PQ}(f)$ may be zero for some path Q that is used in every Wardrop equilibrium. Hence, we also require $\sigma_{PQ}(f) > 0$ for all $P, Q \in \mathcal{P}$ and every f that can be reached from an initial state f_0. Note that the proportional sampling rule satisfies this condition if $f_0 \in \mathrm{int}(\mathcal{F}_\Gamma)$.

Let us now give two examples for the migration rule μ. A natural, but, as we will see, deceptive, migration rule is the *better response policy* defined by

$$\mu(\ell_P, \ell_Q) = \begin{cases} 1 & \text{if } \ell_Q < \ell_P \\ 0 & \text{otherwise} \end{cases} .$$

Another natural migration rule is to switch from P to Q with a probability proportional to the latency difference:

$$\mu(\ell_P, \ell_Q) = \begin{cases} \frac{\ell_P - \ell_Q}{\ell_{\max}} & \text{if } \ell_Q < \ell_P \\ 0 & \text{otherwise} \end{cases} .$$

This corresponds to the migration rule of the replicator dynamics. We call this migration rule *linear migration policy*.

Again, we require that the migration rule is Lipschitz-continuous. All migration rules $\mu(\cdot, \cdot)$ we consider here satisfy $\mu(\ell_P, \ell_Q) > 0$ if and only if $\ell_Q > \ell_P$ and are non-decreasing in ℓ_P and non-increasing in ℓ_Q. By requiring the migration probability to be positive if the sampled path offers a strict improvement we can guarantee convergence towards Wardrop equilibria.

As discussed above we need to assume that the migration rule is not too greedy. We formalise this requirement in the following way:

Definition 3.1 (α-smoothness). *We say that a migration-rule $\mu : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \mapsto [0,1]$ is* α-smooth *for some $\alpha > 0$ if for all $\ell_P \geq \ell_Q \geq 0$ the migration probability satisfies $\mu(\ell_P, \ell_Q) \leq \alpha \cdot (\ell_P - \ell_Q)$.*

Note that if μ can be written as a function of $\ell_P - \ell_Q$ then α-smoothness corresponds to Lipschitz-continuity at 0 with Lipschitz parameter α. The better response migration rule is not α-smooth for any α and therefore is not smooth at all. The linear migration policy is $(1/\ell_{\max})$-smooth. The combination of proportional sampling and linear migration policy corresponds to the replicator dynamics considered in Chapter 2 with $\lambda_i(f) = 1/\ell_{\max}$.

Assuming that each agent's amount of traffic is infinitesimally small we can now describe the evolution of the population shares over time in terms of the fluid limit. Given the sampling and migration rule and normalising the Poisson activation rate of the agents to be 1, we can compute the rate $\rho_{PQ}(f)$ at which agents migrate from path P to path Q:

$$\rho_{PQ}(f) = f_P \cdot \sigma_{PQ}(f) \cdot \mu(\ell_P(f), \ell_Q(f)) \ .$$

Summing up over all paths Q we we obtain a system of ordinary differential equations describing the growth rates (or derivatives with respect to time) of the population share utilising path P:

$$\begin{aligned} \dot{f}_P(t) &= \sum_{Q \in \mathcal{P}_i} (\rho_{QP}(f(t)) - \rho_{PQ}(f(t))) \\ f(0) &= f_0 \ . \end{aligned} \tag{3.1}$$

Due to our definition of $\mu(\cdot,\cdot)$, only one of the terms $\rho_{QP}(f)$ and $\rho_{PQ}(f)$ is positive. If the policy satisfies the properties discussed above, the right-hand-side of this equation is Lipschitz-continuous which, for any initial population f_0, guarantees the existence of a unique solution $\xi(f_0, t)$, $t \geq 0$, to Equation (3.1).

Let us also consider the *best response policy* which has already been mentioned in the introduction (Section 1.2.1.1). This policy is not based on sampling but shortest-path computations. Here, each agent reconsidering her strategy chooses a *best response*, i. e., a shortest path with respect to latency, belonging to her own commodity. Again, denote the set of all best replies by $\beta(f)$. Since the shortest path need not be unique, this leads to a differential inclusion rather than a differential equation.

$$\begin{aligned} \dot{f} &\in \beta(f) - f \\ f(0) &= f_0 \ . \end{aligned} \tag{3.2}$$

Observe that the best response policy can be approximated by a policy in the form of Equation (3.1). To that end, let b_i denote the number of best response paths of commodity i and choose the sampling rate $\sigma_{PQ}(f) = 1/b_i$ for any best response path Q. Furthermore use the migration rule $\mu(\ell_1, \ell_2) = 1$ for $\ell_1 > \ell_2$ and $\mu(\ell_1, \ell_2) = 0$ otherwise. These discontinuous functions can be approximated arbitrarily well by continuous functions such that the overall system approximates the best response policy. One way to do this is, e. g., by using the logit dynamics has been presented in Section 1.2.1.1.

Under stale information we will show that this dynamics does not converge towards an equilibrium even for simple networks with two links only.

3.2.2 Stale Information

The model of stale information we use in this chapter was originally introduced by Mitzenmacher [52]. In this model, agents do not observe latency and flow values in the network directly, but rather consult a bulletin board on which this information is posted. The bulletin board is updated only infrequently, at intervals of length T. Systems that may exhibit similar behaviour include networks in which latency information is gathered centrally and broadcast at regular intervals or uploaded to a server from which it may be polled by the agents whenever needed.

The information on the bulletin board includes at least the path latencies $(\ell_P)_{P\in\mathcal{P}}$ which may be given also implicitly as edge latencies $(\ell_e)_{e\in E}$. We assume that the migration rule $\mu(\cdot,\cdot)$ is based on the values on the bulletin board. In addition, the sampling rule $\sigma_{PQ}(f)$ may depend on the current flow. In this case, we assume that the value of $(f_P)_{P\in\mathcal{P}}$ is available on the bulletin board, too.

Using the information on the bulletin board for the evaluation of the sampling and migration rule, our expression for $\dot{f}(t)$ now depends not only on $f(t)$ but also on $f(\hat{t})$ where $\hat{t}$ marks the beginning of the respective update period. Fix an update period length T and let $\hat{t} = \lfloor t/T \rfloor \cdot T$ denote the beginning of a phase containing t. Then, the migration rate becomes

$$\hat{\rho}_{PQ}(f(t)) = f_P \cdot \sigma_{PQ}(f(\hat{t})) \cdot \mu(\ell_P(f(\hat{t})), \ell_Q(f(\hat{t})))$$

and our dynamic system becomes

$$\begin{aligned} \dot{f}_P(t) &= \sum_{Q\in\mathcal{P}_i} \left(\hat{\rho}_{QP}(f(t)) - \hat{\rho}_{PQ}(f(t))\right) \\ f(0) &= f_0 \ . \end{aligned} \tag{3.3}$$

Note that the right hand side of this differential equation may now have discontinuities whenever the bulletin board is updated. However, since it is Lipschitz-continuous within every phase, there exists a unique solution for each update period and, hence, on all of $\mathbb{R}_{\geq 0}$.

Correspondingly, the best response policy under stale information is then

$$\begin{aligned} \dot{f}(t) &\in \beta(f(\hat{t})) - f(t) \\ f(0) &= f_0 \ . \end{aligned} \tag{3.4}$$

As a convention, whenever the value of t is clear from context, all symbols marked with a hat (ˆ) will refer to time $\hat{t}$, e. g., we may write f instead of $f(t)$ and $\hat{f}$ instead of $f(\hat{t})$. Similarly, we write $\hat{\ell}_e = \ell_e(f(\hat{t}))$, etc.

3.3 Stale Information Makes Convergence Hard

We start by showing that any of our rerouting policies converges towards the set of Wardrop equilibria when information is always up-to-date. On the contrary, we show that the best response policy fails to converge however small the update period is chosen.

3.3.1 Convergence Under Up-To-Date Information

Assuming that information is always up-to-date, convergence towards Wardrop equilibria can easily be shown using the Beckmann-McGuire-Winsten potential [11] Φ as defined in Equation (1.8). Again, Lipschitz-continuity of $\sigma_{PQ}(\cdot)$ and $\mu(\cdot,\cdot)$ ensures existence and uniqueness of a solution to the dynamical system given by Equation (3.1) due to the Picard-Lindelöf-Theorem [16].

Theorem 3.1. *Suppose that $\sigma_{PQ}(f)$ is Lipschitz continuous in f and positive for all $i \in [k]$ and all $P, Q \in \mathcal{P}_i$, and $\mu(\ell_P, \ell_Q)$ is Lipschitz continuous, non-negative, and zero only if $\ell_P \leq \ell_Q$. Then, for any $f_0 \in \mathcal{F}_\Gamma$, the solution $\xi(f_0, t)$ of the dynamical system given by Equation (3.1) converges towards the set of Wardrop equilibria, i. e., $\xi(f_0, t)_{t\to\infty} \to \mathcal{F}_\Gamma^{NE}$.*

Proof. Again, the proof of convergence is an application of Lyapunov's direct method (Theorem 1.7). Consider the derivative with respect to time of Φ along a solution of Equation (3.1).

$$\dot{\Phi}(f) = \sum_{e\in E} \dot{f}_e \cdot \ell_e(f_e) = \sum_{P\in\mathcal{P}} \dot{f}_P \cdot \ell_P(f) = \sum_{P,Q\in\mathcal{P}} \rho_{PQ} \cdot (\ell_Q - \ell_P) \ .$$

By definition of $\mu(\cdot,\cdot)$ and $\sigma_{PQ}(\cdot)$ we know that ρ_{PQ} is positive if $\ell_Q - \ell_P < 0$ and zero otherwise. Hence, $\dot{\Phi}(f) \leq 0$. The inequality is strict if f is not at a Wardrop equilibrium. Consequently, $\Phi(\cdot) - \Phi^*$ is a Lyapunov function and the solution orbit points strictly into every contour set of Φ implying that the solution approaches the set of Wardrop equilibria. $\square$

3.3.2 Oscillation of Best-Response

Similarly, the best response-dynamics converges towards a Wardrop equilibrium under up-to-date information. However, under stale information, the best response dynamics oscillates, no matter how small we choose the update period length T.

We consider an instance with two parallel links with continuous latency functions $\ell_1(x) = \ell_2(x) = \max\{0, \beta \cdot (x - \frac{1}{2})\}$ for a parameter $\beta > 0$ and demand $r = 1$. Note that the latency functions have maximum slope β. At a Wardrop equilibrium we have $f_1 = f_2 = 1/2$ and $\ell_1(f_1) = \ell_2(f_2) = 0$.

First observe that the solution orbit of the best response dynamics (3.4) is unique within any update period and therefore on all of $\mathbb{R}_{\geq 0}$. Consider an update period starting at time t_0 such that $f_1 \neq 1/2 \neq f_2$. Then, for $t \leq T$, the best reply is uniquely defined and the population on the other link decreases exponentially fast, i. e.,

$$f_1(t_0 + t) = \begin{cases} f_1(t_0) \cdot \exp(-t) & \text{if } \ell_1(f_1(t_0)) > \ell_2(f_2(t_0)) \\ 1 - (1 - f_1(t_0)) \cdot \exp(-t) & \text{if } \ell_1(f_1(t_0)) < \ell_2(f_2(t_0)) \end{cases}$$

and $f_2(t) = 1 - f_1(t)$. Note that for our instance, the condition $\ell_1(f_1(t)) > \ell_2(f_2(t))$ is equivalent to $f_1(t) > 1/2$. As an initial condition, we choose

$$f_1(0) = \frac{1}{\exp(-T) + 1} \qquad \text{and} \qquad f_2(0) = \frac{\exp(-T)}{\exp(-T) + 1} \ .$$

Since $f_1(0) > 1/2$ we have $f_1(T) = f_1(0) \cdot \exp(-T) < 1/2$ and

$$\begin{aligned} f_1(2\,T) &= 1 - (1 - f_1(T)) \cdot \exp(-T) \\ &= 1 - \left(1 - \frac{1}{\exp(-T) + 1} \cdot \exp(-T)\right) \cdot \exp(-T) \\ &= f_1(0) \ . \end{aligned}$$

The solution oscillates coming back to its initial configuration every other round. How small do we have to make T in order to guarantee that the maximum latency does not exceed the Wardrop latency (which, in this case, is zero) by more than ϵ? At the beginning of every round i the deviation from the Wardrop latency is

$$\begin{aligned} X &= \max\{\ell_1(f_1(i \cdot T)),\ \ell_2(f_2(i \cdot T))\} \\ &= \beta \cdot \left(\frac{1}{\exp(-T) + 1} - \frac{1}{2}\right) \\ &= \frac{\beta \cdot (1 - \exp(-T))}{2 \cdot \exp(-T) + 2} \ . \end{aligned}$$

Observe that this latency is sustained by more than one half of the agents. In order to guarantee that $X \leq \epsilon$ we must have

$$T \leq \ln\left(\frac{1+2\,\epsilon/\beta}{1-2\,\epsilon/\beta}\right) = \ln\left(1+2\frac{\epsilon}{\beta}\right) - \ln\left(1-2\frac{\epsilon}{\beta}\right) = \mathcal{O}\left(\frac{\epsilon}{\beta}\right) \ .$$

In particular, we can avoid oscillation of the best response policy only for $T = 0$, i. e., in the case of current information.

3.4 Convergence Under Stale Information

We have seen that under up-to-date information the potential decreases. More precisely, if agents migrate from path P to path Q at rate ρ_{PQ}, they cause an infinitesimal potential gain of $\rho_{PQ} \cdot (\ell_Q - \ell_P)$ which is negative for selfish policies. Consider a phase in the bulletin board model starting with a population $\hat{f} \in \mathcal{F}_\Gamma$ and ending with a population $f \in \mathcal{F}_\Gamma$. Let Δf_{PQ} denote the fraction of agents migrating from path P to path Q within this phase. Since throughout the phase the latency values on the bulletin board are fixed to the values they had at the beginning of the phase, agents "see" a *virtual potential gain* of

$$V_{PQ} = \Delta f_{PQ} \cdot (\ell_Q - \ell_P) \ . \tag{3.5}$$

Unfortunately, summing up over all pairs of paths, this does not yield the true potential gain in one round since the latency values do actually change during the phase. The *error terms*

$$U_e = \int_{\hat{f}_e}^{f_e} (\ell_e(u) - \ell_e(\hat{f}_e))\, du \tag{3.6}$$

account for this difference. The following lemma shows that the virtual potential gain terms and the error terms sum up to the true potential gain.

Lemma 3.2. *Let $\hat{f}, f \in \mathcal{F}_\Gamma$ and V_{PQ} and U_e be defined as in Equations (3.5) and (3.6), respectively. Then,*

$$\Phi(f) - \Phi(\hat{f}) = \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} V_{PQ} \ .$$

The virtual potential gain of one round in which the flow vector changes from $\hat{f}$ to f is denoted by

$$\mathcal{V}(\hat{f}, f) = \sum_{P,Q} V_{PQ} = \sum_{e \in E} \ell_e(\hat{f}) \cdot (f_e - \hat{f}_e) \ . \tag{3.7}$$

For the class of policies under consideration, every term V_{PQ} is non-negative and at least one is negative. The following lemma shows that if the rerouting policy is smooth enough, the real potential gain is still at least half of the virtual potential gain and is thus, in particular, negative. Recall that D denotes an upper bound on the length of any path $P \in \mathcal{P}$.

Lemma 3.3. *Suppose that there exist constants $\alpha, \beta > 0$ such that*

1. *the migration rule $\mu(\cdot,\cdot)$ is α-smooth,*
2. *the slope of the latency functions is bounded from above by β, i. e., for all $e \in E$, $u \in [0,1]$, $\ell_e'(u) \leq \beta$.*

Define $T = 1/(4\,D\,\alpha\,\beta)$ and consider a phase beginning at time $\hat{t}$ with an update of the bulletin board and ending at time $t = \hat{t} + \tau$, $\tau \leq T$. For any solution of a routing policy of the form of Equation (3.3) it holds that

$$\Delta\Phi = \Phi(t) - \Phi(\hat{t}) \leq \frac{1}{2} \cdot \mathcal{V}(f(\hat{t}), f(t)) \leq 0 \ .$$

To avoid unnecessary confusion, let us remark that this upper bound on $\Delta\Phi$, which is negative, is actually a lower bound on its absolute value. We can now proof both lemmas.

Proof of Lemma 3.2. We write $\hat{\ell}_e = \ell_e(\hat{f}_e)$. By definition of Φ,

$$\begin{aligned}
\Phi(f) - \Phi(\hat{f}) &= \sum_{e \in E} \int_{\hat{f}_e}^{f_e} \ell_e(u)\,du \\
&= \sum_{e \in E} \left(\int_{\hat{f}_e}^{f_e} (\ell_e(u) - \hat{\ell}_e)\,du + (f_e - \hat{f}_e) \cdot \hat{\ell}_e \right) \\
&= \sum_{e \in E} \left(U_e + \sum_{P \in \mathcal{P}} \sum_{Q \ni e} (\Delta f_{PQ} - \Delta f_{QP}) \cdot \hat{\ell}_e \right) \ .
\end{aligned}$$

Now observe, that the term $\Delta f_{PQ} \cdot \hat{\ell}_e$ occurs positively for all $e \in Q$ and negatively for all $e \in P$. Hence,

$$\begin{aligned}
\Phi(f) - \Phi(\hat{f}) &= \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} \Delta f_{PQ} \left(\sum_{e \in Q} \hat{\ell}_e - \sum_{e' \in P} \hat{\ell}_{e'} \right) \\
&= \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} \Delta f_{PQ} \cdot (\hat{\ell}_Q - \hat{\ell}_P) \\
&= \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} V_{PQ} \ ,
\end{aligned}$$

our claim. □

Proof of Lemma 3.3. We consider a phase starting at time $\hat{t}$ and ending at time $t = \hat{t} + \tau$. For brevity, we will omit the arguments $\hat{t}$ and t, respectively, and write $\hat{f} = f(\hat{t})$ and $f = f(t)$. In the following, a hat (ˆ) indicates symbols referring to values at time $\hat{t}$ whereas symbols without a hat refer to values at time t. The difference between two values between time $\hat{t}$ and t is indicated by a preceding Δ, e. g., $\Delta f_e = f_e - \hat{f}_e$.

By Lemma 3.2, the potential change in one round is

$$\Delta\Phi = \Phi - \hat{\Phi} = \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} V_{PQ} \ . \tag{3.8}$$

We know that the virtual potential gain is negative. The error terms may be positive, but we will see that their absolute value is at most half of the absolute value of the virtual potential gain. To that end, we partition the virtual potential gain terms into chunks

$$V_{PQ}^e = \begin{cases} \frac{V_{PQ}}{4\,D} & \text{if } e \in P \cup Q \\ 0 & \text{otherwise} \ , \end{cases}$$

which we will set against the error terms. Since the paths $P, Q \ni e$ contain at most $2\,D$ edges altogether and since $V_{PQ} \le 0$, we have

$$\sum_{e \in E} V_{PQ}^e \ge 2\,D \cdot \frac{V_{PQ}}{4\,D} = \frac{1}{2} \cdot V_{PQ} \ .$$

Substituting this into Equation (3.8) yields

$$\begin{aligned} \Delta\Phi &\le \sum_{e \in E} U_e + \sum_{P,Q \in \mathcal{P}} \left(\sum_{e \in E} V_{PQ}^e + \frac{1}{2} V_{PQ} \right) \\ &= \sum_{e \in E} \left(U_e + \sum_{P,Q \in \mathcal{P}} V_{PQ}^e \right) + \frac{1}{2} \sum_{P,Q \in \mathcal{P}} V_{PQ} \\ &= \sum_{e \in E} \left(U_e + \sum_{P,Q \in \mathcal{P}} V_{PQ}^e \right) + \frac{1}{2} \cdot \mathcal{V}(\hat{f}, f) \ . \end{aligned}$$

To obtain the assertion of the lemma it remains to show that the first sum is non-positive. We show that this holds termwise, i. e., for all $e \in E$,

$$U_e \le - \sum_{P,Q \in \mathcal{P}} V_{PQ}^e \ . \tag{3.9}$$

We prove this inequality by writing U_e as an upper sum and using an inductive argument on the number of terms in this sum. Fix an edge $e \in E$ and, for the time being, suppose that e is a "growing" edge, i. e., $f_e > \hat{f}_e$. We consider pairs of paths (P_i, Q_i) with $e \in Q_i$ and $\hat{\ell}_{P_i} > \hat{\ell}_{Q_i}$ in ascending

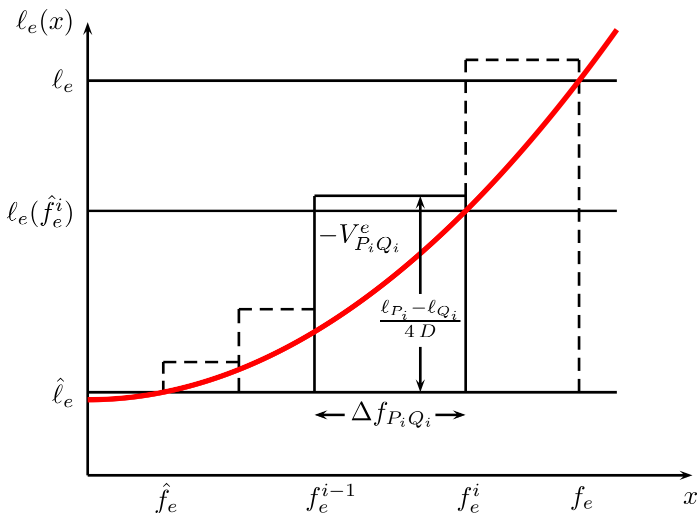

Figure 3.1: Bounding the error term U_e piecewise. The area under the function ℓ_e between $\hat{f}_e$ and f_e is the true contribution to the potential difference. The area between $\hat{\ell}_e$ and $\ell_e(\cdot)$ is the error term U_e. The rectangular areas are the virtual potential gain terms V^e_{PQ} seen by the respective pair of paths P, Q. Each such area contains the respective section of U_e.

order of $\hat{\ell}_{P_i} - \hat{\ell}_{Q_i}$ for $i \in \{1, \ldots, p\}$, $p \leq |\mathcal{P}|^2$ being the number of pairs of paths. Let

$$f_e^0 = \hat{f}_e \qquad \text{and} \qquad f_e^i = f_e^{i-1} + \Delta f_{P_iQ_i} \quad \text{for } i > 0$$

and observe that $f_e \leq f_e^p$. We show that

$$\int_{\hat{f}_e}^{f_e^i} (\ell_e(u) - \hat{\ell}_e)\, du \leq -\sum_{j=1}^{i} V^e_{P_jQ_j}$$

which, for $i = p$, implies Inequality (3.9) and concludes the proof. The situation is depicted in Figure 3.1.

The claim certainly holds for $i = 0$. For the induction step, we derive an upper bound on the growth of the flow and the latency on resource e due to the first i path pairs. By the dynamics given in Equation (3.3), we can upper bound the increase of the flow:

$$\begin{aligned}
f_e^i - \hat{f}_e &= \int_{\hat{t}}^{t} \sum_{j=1}^{i} f_{P_j}(u) \cdot \sigma_{P_jQ_j}(\hat{f}) \cdot \mu(\hat{\ell}_{P_j}, \hat{\ell}_{Q_j})\, du \\
&\leq \int_{\hat{t}}^{t} \sum_{j=1}^{i} f_{P_j}(u) \cdot \sigma_{P_jQ_j}(\hat{f}) \cdot \alpha \cdot (\hat{\ell}_{P_j} - \hat{\ell}_{Q_j})\, du \\
&\leq \alpha \cdot (\hat{\ell}_{P_i} - \hat{\ell}_{Q_i}) \cdot \int_{\hat{t}}^{t} \sum_{j=1}^{i} f_{P_j}(u) \cdot \sigma_{P_jQ_j}(\hat{f})\, du \\
&\leq T \cdot \alpha \cdot (\hat{\ell}_{P_i} - \hat{\ell}_{Q_i}) \ . \qquad (3.10)
\end{aligned}$$

In the first inequality we have used α-smoothness of $\mu(\cdot,\cdot)$. The second inequality is due to the fact that $\hat{\ell}_{P_j} - \hat{\ell}_{Q_j} \leq \hat{\ell}_{P_i} - \hat{\ell}_{Q_i}$ by our ordering of the (P_j, Q_j). For the last step we have used $\sigma_{PQ} \leq 1$, $\sum_{P \in \mathcal{P}} f_P = 1$, and $t - \hat{t} \leq T$. Furthermore, since the slope of ℓ_e is bounded from above by β, using Equation (3.10) we have

$$\begin{aligned}
\ell_e(f_e^i) - \ell_e(\hat{f}_e) &\leq (f_e^i - \hat{f}_e) \cdot \beta \\
&\leq T \cdot \alpha \cdot \beta \cdot (\hat{\ell}_{P_i} - \hat{\ell}_{Q_i}) \ . \qquad (3.11)
\end{aligned}$$

Using this upper bound, we can now prove our claim:

$$\begin{aligned}
\int_{\hat{f}_e}^{f_e^i} (\ell_e(u) - \hat{\ell}_e)\, du &= \int_{\hat{f}_e}^{f_e^{i-1}} (\ell_e(u) - \hat{\ell}_e)\, du + \int_{f_e^{i-1}}^{f_e^i} (\ell_e(u) - \hat{\ell}_e)\, du \\
&\leq -\sum_{j=1}^{i-1} V_{P_jQ_j}^e + (f_e^i - f_e^{i-1}) \cdot (\ell_e(f_e^i) - \hat{\ell}_e) \\
&= -\sum_{j=1}^{i-1} V_{P_jQ_j}^e + \Delta f_{P_iQ_i} \cdot (\ell_e(f_e^i) - \hat{\ell}_e) \\
&\leq -\sum_{j=1}^{i-1} V_{P_jQ_j}^e + \Delta f_{P_iQ_i} \cdot T \cdot \alpha \cdot \beta \cdot (\hat{\ell}_{P_i} - \hat{\ell}_{Q_i})
\end{aligned}$$

where the first inequality uses the induction hypothesis and monotonicity of ℓ_e and the second uses Inequality (3.11). Finally, by definition of $V_{P_iQ_i}^e = (\hat{\ell}_{Q_i} - \hat{\ell}_{P_i}) \cdot \Delta f_{P_iQ_i} / (4\,D)$ and the definition of $T = 1/(4\,D\,\alpha\,\beta)$,

$$\begin{aligned}
\int_{\hat{f}_e}^{f_e^i} (\ell_e(u) - \hat{\ell}_e)\, du &\leq -\sum_{j=1}^{i-1} V_{P_jQ_j}^e - V_{P_iQ_i}^e \cdot 4 \cdot D \cdot T \cdot \alpha \cdot \beta \\
&= -\sum_{j=1}^{i} V_{P_jQ_j}^e \ ,
\end{aligned}$$

which concludes our inductive argument. To see that the same argument holds for edges $e \in E$ with $f_e < \hat{f}_e$, we consider pairs of paths P_i, Q_i with $e \in P_i$. In this case, for Equation (3.10) we obtain $f_e^i - \hat{f}_e \geq T \cdot \alpha \cdot (\hat{\ell}_{Q_i} - \hat{\ell}_{P_i})$. Analogously, for Equation (3.11) we obtain $\ell_e(f_e^i) - \ell_e(\hat{f}_e) \geq T \cdot \alpha \cdot \beta \cdot (\hat{\ell}_{Q_i} - \hat{\ell}_{P_i})$. The remainder of the proof is unchanged. □

Convergence is a direct consequence of this lemma.

Corollary 3.4. *Suppose that there exist constants $\alpha, \beta > 0$ such that*

1. *the migration rule $\mu(\cdot, \cdot)$ is α-smooth,*
2. *the slope of the latency functions is bounded from above by β, i. e., for all $e \in E$, $u \in [0, 1]$, $\ell_e'(u) \leq \beta$,*

and the bulletin board is updated at intervals of length $T \leq 1/(4\,D\,\alpha\,\beta)$. Furthermore, suppose that $\sigma_{PQ}(f)$ is Lipschitz continuous in f and positive for all $i \in [k]$ and all $P, Q \in \mathcal{P}_i$, and $\mu(\ell_P, \ell_Q)$ is Lipschitz continuous, non-negative, and zero only if $\ell_P \leq \ell_Q$. Then, for any $f_0 \in \mathcal{F}_\Gamma$, the solution $\xi(f_0, t)$ of the dynamical system given by Equation (3.3) in the bulletin board model converges towards the set of Wardrop equilibria, i. e., $\xi(f_0, t)_{t \to \infty} \to \mathcal{F}_\Gamma^{NE}$.

Proof. Let $\Phi^* = \min_f \Phi(f)$ denote the optimal potential. It is sufficient to show that for any $\tilde{\Phi} > \Phi^*$ there exists a phase starting with a flow f such that $\Phi(f) \leq \tilde{\Phi}$. Consider the convex set $C_{\tilde{\Phi}}$ containing all points f with $\Phi(f) \leq \tilde{\Phi}$. Note that $C_{\tilde{\Phi}}$ contains all Wardrop equilibria. Let $\mathcal{V}(f)$ denote the virtual potential gain in one update period starting with population f. Note that $\mathcal{V}(\cdot)$ is non-positive, continuous in f, and zero only if f is a Wardrop equilibrium. Furthermore, $C_{\tilde{\Phi}}$ is compact. Hence, there exists an $\epsilon_{\tilde{\Phi}} > 0$ such that

$$\inf_{f \notin C_{\tilde{\Phi}}} \{-\mathcal{V}(f)\} \geq \epsilon_{\tilde{\Phi}} \ .$$

Let $\Delta\Phi(f)$ denote the true potential gain of an update period starting with population f. Lemma 3.3 asserts that

$$\inf_{f \notin C_{\tilde{\Phi}}} \{-\Delta\Phi(f)\} \geq \frac{\epsilon_{\tilde{\Phi}}}{2} \ .$$

Now, consider a sequence of flows $(f(i\,T))_{i \in \{0,\ldots,k\}}$ with $f(i\,T) \notin C_{\tilde{\Phi}}$. Then, the potential is at most

$$\Phi(f(k\,T)) \leq \Phi(f(0)) - k \cdot \frac{\epsilon_{\tilde{\Phi}}}{2} \ ,$$

and, since $\Phi(f(k\,T)) \geq \tilde{\Phi}$ by assumption, any such sequence must be of finite length. □

3.5 Convergence Time

Although we have seen that our smooth adaptive rerouting policies converge towards Wardrop equilibria in the long run, none of them will necessarily reach an exact equilibrium. In order to give bounds on the convergence time we consider approximate equilibria similar to the ones defined in Chapter 2. We consider both the uniform and the proportional sampling policy in combination with the linear migration policy. Recall that for the linear migration policy $\mu(\ell_P, \ell_Q) = \max\{(\ell_P - \ell_Q)/\ell_{\max}, 0\}$ we have $\alpha = 1/\ell_{\max}$. Again, let β be an upper bound on the slope of the latency functions.

Our definition of approximate equilibria differs slightly from the one used in Chapter 2. As for our definition of ϵ-equilibria (Definition 2.2), we again require almost all of the agents to sustain latency that is not much greater than the latency sustained by other agents of the same commodity. Note that for the linear migration rule the migration probability is proportional to the *absolute* latency difference between two paths rather than the *ratio* between the difference and the average latency as was the case for the replicator dynamics considered in Chapter 2. This is reflected by our definition of approximate equilibria in that it allows for an additive (rather than a relative) deviation.

Our analyses will disclose a fundamental difference between the proportional and the uniform sampling rule. We will see that for the uniform sampling rule we can prove convergence towards a stricter definition of approximate equilibria at the cost of a factor $|\mathcal{P}|$ in the time of convergence.

3.5.1 Uniform Sampling

Recall that the uniform sampling rule is defined as $\sigma_{PQ} = 1/|\mathcal{P}_i|$ for $P, Q \in \mathcal{P}_i$ and any commodity $i \in [k]$.

Definition 3.2 (additive (δ, ϵ)-equilibrium)**.** *For a flow $f \in \mathcal{F}_\Gamma$, let $\ell^i_{\min} = \min\{\ell_P(f) \mid P \in \mathcal{P}_i\}$ denote the minimum latency of commodity $i \in [k]$. For every commodity i and every path $P \in \mathcal{P}_i$ the agents on path P are said to be* δ-unsatisfied *iff $\ell_P(f) > \ell^i_{\min} + \delta$. The flow f is said to be at an* additive (δ, ϵ)-equilibrium *iff the total volume of δ-unsatisfied agents is at most ϵ.*

For the uniform sampling policy the growth rate of the flow on every path is bounded by the inverse of the number of paths. Hence, our time bounds must also depend on this parameter.

Theorem 3.5. *Suppose that the bulletin board is updated at intervals of length*

$$T \leq \frac{1}{4\, D\, \alpha\, \beta} \enspace .$$

For the uniform sampling policy combined with the linear migration policy, the number of update periods not starting at an additive (δ, ϵ)-equilibrium is bounded from above by

$$\mathcal{O}\left(\frac{\max_{i\in[k]}\{|\mathcal{P}_i|\}}{\epsilon \cdot T} \cdot \left(\frac{\ell_{\max}}{\delta}\right)^2\right) \quad .$$

Proof. Let $m = \max_{i\in[k]}\{|\mathcal{P}_i|\}$. We consider a phase that does not start at an additive (δ, ϵ)-equilibrium and compute the virtual potential gain of the phase. First note that the population shares cannot decrease much within one round. Since for every $P \in \mathcal{P}$, $\dot{f}_P \geq -f_P$, at the end of the update period $f_P(t+T) \geq f_P(t) \cdot \exp(-T)$. In particular, we know that the volume of δ-unsatisfied agents is at least ϵ at the beginning of the phase and hence at least $\epsilon \cdot \exp(-1)$ throughout the phase if $T \leq 1$.

We now consider the total contribution of δ-unsatisfied agents to the virtual potential gain. Let $\mathcal{P}_i^\delta$ denote the set of expensive paths of commodity i, i. e., $\mathcal{P}_i^\delta = \{P \in \mathcal{P}_i \mid \ell_P(t) > \ell^i_{\min} + \delta\}$ and let $\mathcal{P}^\delta = \bigcup_{i\in[k]} \mathcal{P}_i^\delta$. Furthermore, let P_i^* denote a path with minimal latency in commodity i. First note that the contributions of all agents to the virtual potential gain is non-positive since agents only migrate if this decreases their latency. From this fact and Lemma 3.3 it follows that the true potential gain in one round is at least

$$\Delta\Phi \leq -\frac{1}{2} \sum_{P\in\mathcal{P}_\delta} V_{PP_i^*} \quad .$$

For every commodity $i \in [k]$ and every path $P \in \mathcal{P}_i$ the probability for an agent on path P to sample path P_i^* is at least

$$\sigma_{PP_i^*} = \frac{1}{|\mathcal{P}_i|} \geq \frac{1}{m} \quad .$$

If $P \in \mathcal{P}_i^\delta$, then the migration probability is

$$\mu(\ell_P, \ell^i_{\min}) = \frac{\ell_P - \ell^i_{\min}}{\ell_{\max}} \geq \frac{\delta}{\ell_{\max}} \quad .$$

The total virtual potential gain over a phase of length T is hence at least

$$\begin{aligned}
\sum_{P\in\mathcal{P}^\delta} V_{PP_i^*} &\leq T \sum_{P\in\mathcal{P}^\delta} \inf_{u\in[t,t+T)} f_P(u) \cdot \frac{\delta \cdot (\ell_{P_i^*}(t) - \ell_P(t))}{m \cdot \ell_{\max}} \\
&\leq -T \sum_{P\in\mathcal{P}^\delta} f_P(t) \cdot \exp(-1) \cdot \frac{\delta^2}{m \cdot \ell_{\max}} \\
&\leq -T \frac{\epsilon \cdot \delta^2 \cdot \exp(-1)}{m \cdot \ell_{\max}}
\end{aligned}$$

where the last inequality holds because the flow is not at an additive (δ, ϵ)-equilibrium at time t and hence $\sum_{P \in \mathcal{P}^\delta} f_P(t) \geq \epsilon$. By Lemma 3.3, the potential decreases by at least half of this in every round, i. e.,

$$\Delta\Phi \leq -T \frac{\epsilon \cdot \delta^2 \cdot \exp(-1)}{2 \cdot m \cdot \ell_{\max}}$$

as long as we are not at an additive (δ, ϵ)-equilibrium. Suppose that we are not at an additive (δ, ϵ)-equilibrium for n rounds. Then, since $\Phi(f) \leq \ell_{\max}$ for every $f \in \mathcal{F}_\Gamma$ and in particular for $f = f_0$, after these n rounds it holds that

$$\Phi \leq \ell_{\max} - n \cdot T \frac{\epsilon \cdot \delta^2 \cdot \exp(-1)}{2 \cdot m \cdot \ell_{\max}} \quad .$$

Since furthermore $\Phi \geq 0$, it follows that

$$n \leq \mathcal{O}\left(\frac{m}{T \cdot \epsilon}\left(\frac{\ell_{\max}}{\delta}\right)^2\right) \quad ,$$

which is our claim. □

3.5.2 Proportional Sampling

Using proportional instead of uniform sampling we can get rid of the factor $\max_{i \in [k]}\{|\mathcal{P}_i|\}$ at the cost of a slightly weaker definition of approximate equilibria. Recall that the proportional sampling rule is defined as $\sigma_{PQ}(f) = f_Q / r_i$ for all paths $P, Q \in \mathcal{P}_i$ and any commodity $i \in [k]$.

Definition 3.3 (weak additive (δ, ϵ)-equilibrium)**.** *Consider a flow $f \in \mathcal{F}_\Gamma$. For every commodity $i \in [k]$ and every path $P \in \mathcal{P}_i$ the agents on path P are said to be* weakly δ-unsatisfied *iff $\ell_P(f) > L_i(f) + \delta$. The flow f is said to be at a* weak additive (δ, ϵ)-equilibrium *iff the total volume of weakly δ-unsatisfied agents is at most ϵ.*

Note that every additive (δ, ϵ)-equilibrium is also a weak additive (δ, ϵ)-equilibrium.

Theorem 3.6. *Suppose that the bulletin board is updated at intervals of length*

$$T \leq \frac{1}{4\, D\, \alpha\, \beta} \quad .$$

For the proportional sampling policy combined with the linear migration policy, the number of update periods not starting at an additive weak (δ, ϵ)-equilibrium is bounded from above by

$$\mathcal{O}\left(\frac{1}{\epsilon \cdot T} \cdot \left(\frac{\ell_{\max}}{\delta}\right)^2\right) \quad .$$

Proof. We use a similar argument here. Again, consider a phase starting at time t such that $f(t)$ is not at a weak additive (δ, ϵ)-equilibrium. Define $\mathcal{P}_i^\delta = \{P \in \mathcal{P}_i \mid \ell_P(t) > L_i(t) + \delta\}$ and let $\mathcal{P}^\delta = \bigcup_{i \in [k]} \mathcal{P}_i^\delta$. Consider the potential gain of agents using paths in $\mathcal{P}^\delta$ throughout a phase of length T. By Lemma 3.3,

$$\begin{aligned}
\Delta\Phi &\leq \frac{1}{2} \sum_{P,Q \in \mathcal{P}} V_{PQ} \\
&\leq \frac{\exp(-1) \cdot T}{2} \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i^\delta} \sum_{Q \in \mathcal{P}_i} \frac{f_P(t) \cdot f_Q(t) \cdot (\ell_P - \ell_Q)}{r_i \cdot \ell_{\max}} \cdot (\ell_Q - \ell_P) \\
&= -\frac{\exp(-1) \cdot T}{2 \cdot \ell_{\max}} \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i^\delta} f_P(t) \sum_{Q \in \mathcal{P}_i} \frac{f_P(t) \cdot f_Q(t) \cdot (\ell_P - \ell_Q)^2}{r_i} \quad .
\end{aligned}$$

Using Jensen's inequality,

$$\begin{aligned}
\Delta\Phi &\leq -\frac{\exp(-1) \cdot T}{2 \cdot \ell_{\max}} \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i^\delta} f_P(t) \left(\sum_{Q \in \mathcal{P}_i} \frac{f_Q(t) \cdot (\ell_P - \ell_Q)}{r_i} \right)^2 \\
&= -\frac{\exp(-1) \cdot T}{2 \cdot \ell_{\max}} \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i^\delta} f_P(t) (\ell_P - L_i(f(t)))^2 \\
&\leq -\frac{\exp(-1) \cdot \epsilon \cdot \delta^2 \cdot T}{2 \cdot \ell_{\max}} \quad ,
\end{aligned}$$

where for the equality we have used $\sum_{P \in \mathcal{P}_i} f_Q = r_i$ and the definition of L_i, and for the last inequality we have used $\sum_{P \in \mathcal{P}^\delta} f_P \geq \epsilon$. Given a bound on the potential gain per round, the remainder of the proof is analogous to the proof of Theorem 3.5. □

3.6 Summary

We have analysed selfish rerouting policies in a model of time that is more realistic than the one used in Chapter 2 by taking into account the fact that information may be stale. Our analyses use a bulletin board model in which information is updated at regular intervals. We have introduced the class of α-smooth adaption policies that can cope with the issue of stale information if parametrised correctly. For any class of latency functions with bounded slope and any non-zero update interval length we have shown that for some α small enough, α-smooth adaption policies converge towards the set of Wardrop equilibria.

These positive results complement earlier negative ones in similar models showing that stale information may lead to oscillation which in turn decreases performance. The fact that these analyses yield negative results, however, is a consequence of the greediness of the considered policies. E. g., Mitzenmacher [52] considers a policy in which each agent assigns its job to a machine with minimum load. Negative results for greedy policies are confirmed by our analyses. We have seen that the best response dynamics oscillates in our model as well.

For two representatives of our class of smooth adaption policies we have derived bounds on the time to reach approximate equilibria in a bicriterial sense similar to the one used in Chapter 2. The convergence speed of smooth adaption policies is coupled linearly to the smoothness parameter α. Our results on the convergence time reflect this fact in two aspects. First, since the smoothness parameter α must be chosen proportionally to the reciprocal of the maximum slope and the length of the update period, our time bounds also depend on these parameters. Thus, they have a pseudopolynomial character.

Second, our smoothness condition requires the migration probability to be small for small latency improvements, even if this improvement is huge in relation to the latency of the originating path. For this reason, the progress made by our policies is large only if agents make large absolute improvements. Correspondingly, we have to define approximate equilibria in terms of additive deviations as opposed to the relative deviations allowed for approximate equilibria in the sense of Chapter 2.

Whereas the dependence of the convergence speed on the length of the update period is evidently unavoidable, the dependence on the maximum slope is a consequence of the generality of the framework of this chapter. The next chapter will reveal how this dependence can be avoided if one considers a more specific class of policies.

Chapter 4

Fast Convergence by Exploration and Replication

4.1 Introduction

We have seen that convergence towards Wardrop equilibria is, in principle, possible even if information used by the agents is stale. However, our analysis is not yet satisfactory with respect to the speed of convergence. We have slowed down the dynamics by a factor that is proportional to the maximum slope of the latency functions. Unfortunately, this quantity is unbounded for some natural classes of latency functions, e. g., for affine functions or polynomials, as well as for any class that is closed under scaling. Even for any fixed instance the maximum slope may be large. Let us evaluate the convergence time from an algorithmic point of view. We consider the rerouting policy as a distributed algorithm and an instance Γ as its input. Then we can interpret the time to reach approximate equilibria as a function of the encoding length of Γ. Such an encoding encompasses also a representation of the latency functions. For example, a natural representation of polynomial functions is the coefficient representation which lists the coefficients of all monomials. Any bound that is polynomial in the maximum slope (as is the case in Theorems 3.5 and 3.6) is then pseudopolynomial in the representation length of Γ.

The goal of this chapter is to derive bounds on the time of convergence that are polynomial in the representation length of Γ and are thus much stronger than the ones presented in the preceding chapter. To that end, we analyse a very specific rerouting policy, rather than the large class of smooth adaption policies considered in Chapter 3. Why does our proof of the convergence theorem (Lemma 3.3) require to slow down the dynamics by the reciprocal of the maximum slope? The reason for this is a pessimistic assumption we had to make for bounding the volume of agents migrating towards a particular edge. We have assumed that *all* agents might sample a particular edge. In this chapter, we will introduce a rerouting policy that partially resembles the replicator dynamics from Chapter 2 in that it samples agents rather than paths. For such a policy, the probability to sample an edge is bounded by the flow on this edge itself. Consequently, the growth rate of an edge flow is always relative to the current flow. Analyses depending on the maximum slope of the latency functions cannot take this into account. A more appropriate parameter precisely capturing the relative increase of a function value due to a relative increase of its argument is its *elasticity*. In contrast to the maximum slope, this is a very good-natured parameter which, for instance, coincides with the degree of positive polynomials. In this chapter, we identify elasticity as the crucial parameter determining the speed of convergence rather than the maximum slope. Thus, we obtain bounds on the convergence time that are polynomial rather than pseudopolynomial.

4.1.1 Our Results

We study policies consisting of a sampling step and a migration step. As a first approach we consider a policy resembling the replicator dynamics. In the sampling step an agent currently using path P chooses another agent at random and then migrates to this agent's path Q with a probability proportional to the *relative* latency improvement $(\ell_P - \ell_Q)/\ell_P$. We show that if the activation rate of the agents is bounded by the reciprocal of the elasticity of the latency functions, then this policy converges towards the set of Wardrop equilibria.

To analyse the time of convergence we consider a bicriterial definition of approximate equilibria which now allows for a deviation from L_i that is relative to L. More precisely, we require all but an ϵ-fraction of the agents to deviate by no more than δL in both directions from the average latency of their commodity. In the single-commodity case, this can be interpreted as an allocation where almost all agents are almost satisfied. This improves upon the definition of approximate equilibria used in Chapter 3 which allowed for additive deviations which could be by any factor larger than the average latency. We show that the time to reach an approximate equilibrium in this relative sense is bounded from above by

$$\mathcal{O}\left(\frac{d}{\delta^2\,\epsilon^2}\log\left(\frac{\Phi(f_0)}{\Phi^*}\right)\right)$$

where d is the elasticity of the latency functions. Remarkably, this bound is almost independent of the size of the network. It is proven in Section 4.4.1

for the symmetric case (in which it is actually by a factor ϵ smaller) and in Section 4.5 for the asymmetric case.

As has been discussed in Chapter 2 already, bicriterial approximate equilibria are transient. Being at a bicriterial approximate equilibrium, a minority of agents utilising low latency paths may grow to a volume of more than $\epsilon \cdot r$ at a later time, and it may take an arbitrarily long time until this event happens if the initial flow on these paths is arbitrarily small. Furthermore, due to the lack of innovation of the replication policy, a population converges to a Wardrop equilibrium only if all paths are initially used.

For symmetric instances, both of these shortcomings are resolved by changing the sampling mechanism. Instead of sampling another agent at random, we can also sample another path uniformly at random. Thus, we obtain a non-zero lower bound on the flow of any path with below average latency after a single round already. The uniform sampling rule, however, comes with another two, even more significant disadvantages. First, since under this policy the growth rates are not relative to the prevailing flow, the activation rate must be chosen proportionally to the reciprocal of the maximum slope $\ell'_{\max}$, which is exactly what we want to avoid. Second, the probability to sample a particular path is only $1/|\mathcal{P}_i|$, a number which may be exponentially small in the number of edges. Altogether, the uniform sampling rule alone yields only poor upper bounds on the time of convergence. For this reason, we combine both sampling rules preserving the good properties of both. Our hybrid *exploration-replication policy* uses uniform sampling with small probability to explore the strategy space and to guarantee a minimum flow on all paths with below average latency and uses proportional sampling to boost these paths thereafter. Due to the exponential growth caused by the proportional sampling rule, the terms $|\mathcal{P}_i|$ and $\ell'_{\max}$ caused by the uniform sampling rule appear only logarithmically in our time bounds. Altogether, this enables us to obtain polynomial bounds on the time of convergence towards approximations of the Wardrop equilibrium in terms of the Beckmann-McGuire-Winsten potential in the single-commodity case. More precisely, the time to reach a flow with potential at most $(1+\epsilon)\,\Phi^*$ is bounded from above by

$$\operatorname{poly}\left(\frac{d \cdot D}{\epsilon}\right) \cdot \operatorname{polylog}\left(|E| \cdot \frac{\ell'_{\max}}{\ell_{\min}}\right) \cdot \log\left(\frac{\Phi(f_0)}{\Phi^*}\right)$$

where $\ell'_{\max}$ is an upper bound on the maximum slope of the latency functions and $\ell_{\min}$ is the minimum latency of an edge. This bound, which is proven in Section 4.4.2, is polynomial in the representation length of the instance Γ if, e. g., the latency functions are positive polynomials of arbitrary degree in coefficient representation. Thus, the exploration-replication policy behaves like a fully polynomial time approximation scheme (FPAS).

In the technical part of this chapter, both of the above bounds will be slightly refined. First, $\ell'_{\max}$ needs to be an upper bound on the maximum slope in a region close to 0 only. Second, a relative approximation of Φ^* is only possible if $\Phi^* > 0$. Furthermore, our time bounds are finite only if $\Phi^* > 0$ and $\ell_{\min} > 0$. In the case that $\Phi^* = 0$ or $\ell_{\min} = 0$, we can still approximate Φ^* up to a small additive constant α which appears only logarithmically in the time bounds.

The two take home message of this chapter are that first, proportional sampling can significantly improve convergence time and second, the critical parameter that limits the convergence speed is the elasticity of the latency functions. This is confirmed by two lower bounds, which we present in Section 4.6.

For our first lower bound we consider non-oscillating policies that ensure monotonicity with respect to the potential over a basic class of latency functions with elasticity at most d. We show that any rerouting policy of this kind needs time $\Omega(d/\epsilon)$ to approximate the optimal potential within a factor of $1 + \epsilon$. The network underlying this analysis consists of two parallel links only.

Finally, we show that proportional sampling is actually necessary. To that end, we consider policies using fixed sampling properties that do not rely on any pre-knowledge about the latency functions. We show that for any such policy there exists an instance with exponentially many paths such that the time of convergence is proportional to the number of paths and therefore exponential in the network size.

4.1.2 Related Work

The problem of overshooting is also known in convex optimisation. Various heuristics like the distributed algorithm presented by Bertsekas and Tsitsiklis [13] for the non-linear multi-commodity flow problem proceed by shifting certain amounts of flow from one path to another. This amount of flow depends in a linear fashion on the reciprocal of the second derivative of the latency functions. Using marginal cost latency functions, this algorithm can also be applied to compute Wardrop equilibria in a distributed way. In this case, the slowdown is again linear in the first derivative. Recently, Awerbuch *et al.* [7] presented an efficient distributed algorithm for multi-commodity flow problems which is restricted to instances with linear latency functions.

The approach of no-regret routing addresses our problem from the perspective of on-line learning. The *regret* is defined as the difference between an agent's average latency over time and the latency of the best path in hindsight (see, e. g., [8, 9, 15, 27]). The goal is to minimise the regret over time. The bounds obtained here also depend in a pseudopolynomial fashion on network parameters.

Going to discrete models, analyses become more involved since probabilistic effects have to be considered that can be neglected in fluid limit models. Consequently, results obtained for discrete models are more restrictive in other aspects, e. g., by limiting the strategy space to parallel links. Even-Dar and Mansour [28] study distributed and concurrent rerouting policies in such a model. They consider parallel links with speeds. Upper bounds are presented for the case of agents with identical weights. Their algorithms use static sampling rules that explicitly take into account the speeds of the individual links. Berenbrink *et al.* [12] present an efficient distributed protocol for balancing identical jobs on identical machines.

In Section 1.3.4 we have presented several approaches for computing equilibria in discrete and continuous congestion games. In contrast to our work, these are centralised algorithms whereas we analyse how agents can compute or learn an equilibrium in a distributed fashion.

It is known that systems with an infinite number of agents are very powerful. For example, in [5, 58] quadratic dynamic systems are considered in which individuals are mated at random to produce two individuals as offspring. Systems of this kind are able to solve PSPACE-complete problems. In particular, they can compute Wardrop equilibria. This approach however, does not yield a distributed algorithm, since individuals do not have a natural interpretation as participants in a network routing game.

In evolutionary game theory one also considers dynamics which resemble the replicator dynamics but also use a small amount of random mutation [38]. This has an effect similar to the one introduced by our uniform sampling component.

4.1.3 Preliminaries

We consider the Wardrop model introduced in Section 1.3. Again, we normalise the total flow demand $r = 1$ for simplicity.

4.1.3.1 Elasticity

We have argued that a rerouting policy cannot guarantee convergence to Wardrop equilibria if the latency functions make arbitrarily large leaps due to minor shifts of the flow. To restrict the number of agents migrating simultaneously, it must have some information about the sensitivity of the latency functions to small shifts of traffic. Our analysis shows that the following parameter is relevant.

Definition 4.1 (elasticity). *For any positive differentiable function $\ell : D \mapsto \mathbb{R}$, the elasticity of $\ell(\cdot)$ at $x \in D$ is defined as*

$$d(x) = \frac{x \cdot \ell'(x)}{\ell(x)} \ .$$

We say that a function $\ell(\cdot)$ has elasticity d if $d(x) \leq d$ for every $x \in D$ and a class of functions $\mathcal{L}$ has elasticity d if every $\ell \in \mathcal{L}$ has elasticity d.

Intuitively, elasticity can be interpreted as follows. As a very coarse estimate of the slope of a positive continuous function $f(\cdot)$ at x we can use the slope of the secant through the origin and the point $(x, f(x))$. Then, the elasticity of $f(\cdot)$ at x is the factor by which this estimate is wrong. The elasticity of a function at x describes the proportional growth of the function value as a result of a proportional growth of its argument. In economics, this parameter can be used to describe characteristic properties of a market, e. g., the price elasticity of demand or supply.

As an example, the polynomial function $f(x) = a\,x^d$ has elasticity d over the entire range. The exponential function $f(x) = a \cdot \exp(\lambda\, x)$ has elasticity at most λ for $x \in [0, 1]$ reaching its maximum λ at $x = 1$. Delay functions that arise in the theory of M/M/1 queues [54] like $\ell_1(x) = 1/(1 - x)$ and $\ell_2(x) = x/(1 - x)$ do not have bounded elasticity since they have a singularity at 1. The elasticities of ℓ_1 and ℓ_2 are given by $d_1(x) = x/(1 - x)$ and $d_2(x) = 1/(1 - x)$, respectively. If x is separated from 1 by some amount, the elasticity is finite. Figure 4.1 shows the value of the elasticity of both functions. We see that as long as the load is at most 90 % of the capacity, the elasticity of both functions is still at most 10.

The following fact which will be used frequently shows that for latency functions with bounded elasticity small relative changes of the flow cause only small changes of the latency.

Fact 4.1. *For any non-negative, non-decreasing function ℓ with elasticity d it holds that for any $\delta \in [0, 1/(2\,d)]$, that*

$$\ell((1 + \delta) \cdot x) \leq (1 + 2\,d\,\delta) \cdot \ell(x) \ .$$

Proof. Let $0 \leq \delta \leq 1/(2\,d)$. The derivative $\ell'(y)$ of the function ℓ at any point $y \in [x, (1 + \delta)x]$ is at most $d \cdot \ell(y)/y \leq d \cdot \ell((1 + \delta)x)/x$. This yields

$$\begin{aligned} \ell((1 + \delta)\,x) &= \ell(x) + \int_x^{(1+\delta)x} \ell'(u)\,du \\ &\leq \ell(x) + \delta\,x\,\frac{d \cdot \ell((1 + \delta)\,x)}{x} \ . \end{aligned}$$

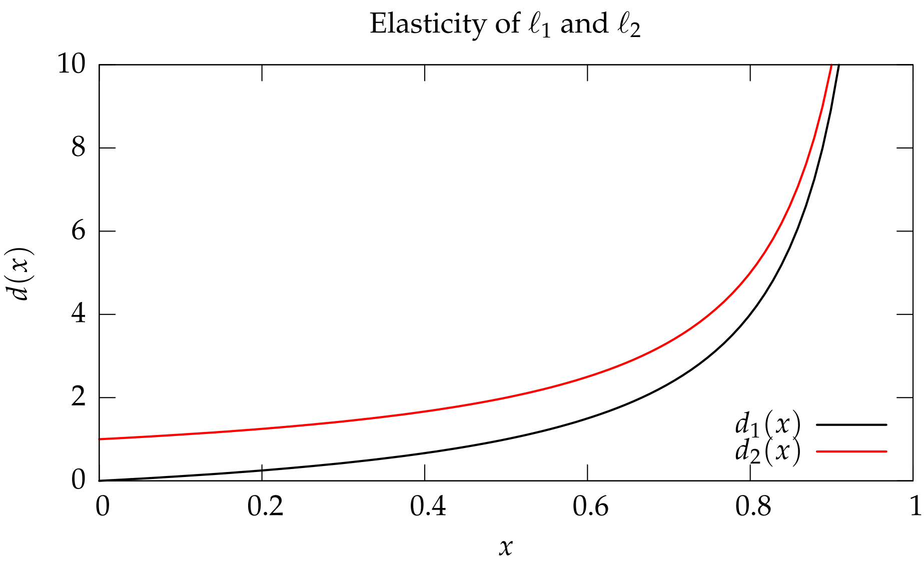

Figure 4.1: The elasticity of two latency functions $\ell_1(x) = 1/(1-x)$ and $\ell_2(x) = x/(1-x)$ which arise in queueing theory.

Hence, $\ell((1+\delta)x) \leq \frac{1}{1-\delta d} \cdot \ell(x)$, and for $\delta \leq 1/(2d)$ we have $\ell((1+\delta)x) \leq (1+2\delta d) \cdot \ell(x)$. $\square$

Let us remark that we can easily generalise the definition of elasticity to functions that are not differentiable by requiring $\ell(x(1+\delta)) \leq (1+\mathcal{O}(d\,\delta))\,\ell(x)$ for $x \in [0,1]$ and $\delta = \mathcal{O}(1/d)$ in the first place.

Moreover, for functions with bounded elasticity, we can bound potential and average latency in terms of each other as the following fact shows.

Fact 4.2. *Let $d \geq 1$ denote an upper bound on the elasticity of the latency functions. Then, for every flow f, $L(f)/(d+1) \leq \Phi(f) \leq L(f)$.*

This can be proved by comparing the contributions of the individual edges to the average and the potential termwise. For any edge $e \in E$, the upper bound follows from the fact that $\int_0^{f_e} \ell_e(u)\,du \leq \ell_e(f_e) \cdot f_e$ by monotonicity of ℓ_e. For the lower bound note that the ratio

$$\frac{\ell_e(f_e) \cdot f_e}{\int_0^{f_e} \ell_e(u)\,du}$$

is maximised if ℓ_e has elasticity over the entire interval $[0,1]$, i.e., $\ell'_e(x) = d \cdot \ell_e(x)/x$ for all $x \in [0,1]$. This functional equation is solved by functions of the form $\ell_e(x) = a \cdot x^d$ for some $a > 0$. For any such a, the above ratio is precisely $d+1$ which is our desired bound. A formal proof can be found in the appendix.

4.1.3.2 α-Shifted Potential

Our goal is the design of distributed rerouting policies that approximate a Wardrop equilibrium. Since Wardrop equilibria minimise the Beckmann-McGuire-Winsten potential [11]

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} \ell(u)\,du \quad ,$$

a natural measure for the quality of approximation, is the ratio between the potential achieved after a certain amount of time and the minimal potential Φ^*. Furthermore, our proofs of the convergence time proceed by showing that the potential Φ decreases by a factor in every round. Consequently our time bounds are logarithmic in the ratio $\Phi(f_0)/\Phi^*$. Note, however, that for the given instance it might be that $\Phi^* = 0$. In this case, none of the above is meaningful, and we suggest to shift the potential by some positive additive term. In general, we consider an *α-shifted potential* of the form $\Phi + \alpha$ where $\alpha \geq 0$ can be chosen arbitrarily in such a way that, for the given instance, $\Phi^* + \alpha$ is strictly positive. The shift of the potential can be interpreted as appending a virtual resource with constant latency α to every path, thus increasing the latency observed by the agents by α.

Definition 4.2 ($\Gamma^{+\alpha}$). *For any instance Γ with potential function Φ let $\Gamma^{+\alpha}$ denote an instance obtained from Γ by appending a new resource e_P to every strategy $P \in \mathcal{P}$ with constant latency function $\ell_{e_P}(x) = \alpha$ and let $\Phi^{+\alpha}$ denote the respective potential function.*

Our policies make use of the parameter α by adding this virtual latency offset to every path. This will be made precise by Fact 4.3 below.

4.2 The Exploration-Replication Policy

In the following, we formally introduce our rerouting policy for any class of latency functions for which the elasticity is bounded by $d \geq 1$. We consider the policy in a round-based version of the Wardrop model resembling the bulletin board model. Whereas in Chapter 3 we have assumed that migration happens throughout an entire update period based on the information posted on the bulletin board, we now assume for simplicity, that agents act in a round based fashion, i. e., within a round they make a single decision only. Thus, we normalise the update period length to be $T = 1$.

Apart from the elasticity d, our policy depends on two parameters α and β. In every round, an agent is activated with constant probability λ/d where $\lambda = 1/32$. It then performs the following two steps. Consider an agent in commodity $i \in [k]$ currently utilising path $P \in \mathcal{P}_i$.

1. *Sampling:* With probability $(1-\beta)$ perform step 1(a) and with probability β perform step 1(b).
   (a) *Proportional sampling:* Sample a path $Q \in \mathcal{P}_i$ with probability f_Q/r_i.
   (b) *Uniform sampling:* Sample a path $Q \in \mathcal{P}_i$ with probability $1/|\mathcal{P}_i|$.
2. *Migration:* If $\ell_Q < \ell_P$, migrate to path Q with probability $\frac{\ell_P - \ell_Q}{\ell_P + \alpha}$.

The activation probability λ/d slows down the policy proportionally to the reciprocal of the elasticity of the latency functions. The migration probability $(\ell_P - \ell_Q)/(\ell_P + \alpha)$ is precisely the relative latency improvement (w. r. t. latency functions with a virtual offset α) when migrating from path P to path Q. Whereas the parameter $\alpha \geq 0$ can be chosen arbitrarily, the parameter β must be chosen subject to the constraint

$$\beta \leq \frac{\min_{P\in\mathcal{P}} \ell_P(0) + \alpha}{D \cdot \max_{e\in E} \max_{x\in[0,\beta]} \ell'_e(x)} \ . \tag{4.1}$$

This is the ratio between minimum path latency and the maximum derivative of the path latency for close to zero flow values. In the following, we always assume that this constraint is satisfied. We now define our policy formally in terms of the amount of flow that is shifted between any pair of paths within one round.

Definition 4.3 (exploration-replication policy)**.** *For an instance Γ let $d \geq 1$ be an upper bound on the elasticity of the latency functions and let β be chosen as in Equation (4.1). Consider a round starting with flow f. For every commodity $i \in [k]$ and every path $P, Q \in \mathcal{P}_i$ with $\ell_Q \leq \ell_P$, the (α, β)-*exploration-replication policy *shifts a flow volume of*

$$\rho_{PQ}(f) = \lambda \cdot \frac{1}{d} \cdot f_P \cdot \left((1-\beta) \cdot \frac{f_Q}{r_i} + \beta \cdot \frac{1}{|\mathcal{P}_i|} \right) \frac{\ell_P - \ell_Q}{\ell_P + \alpha} \tag{4.2}$$

agents from path P to path Q, where $\lambda = \frac{1}{32}$.

Given an initial population $f_0 \in \mathcal{F}_\Gamma$, the sequence of flow vectors generated by the (α,β)-exploration-replication policy is given by

$$\begin{aligned} f(0) &= f_0 \\ f_P(t+1) &= f_P(t) - \sum_{Q\in\mathcal{P}_{i_P}} \rho_{PQ}(f(t)) + \sum_{Q\in\mathcal{P}_{i_P}} \rho_{QP}(f(t)) \ . \end{aligned}$$

In our proofs we will simulate the (α, β)-exploration-replication policy by applying the $(0, \beta)$-exploration-replication policy to a modified instance with additional offsets α added to the path latencies. The following fact asserts that this is equivalent.

Fact 4.3. *For some $\alpha \geq 0$ consider an instance Γ and the instance $\Gamma^{+\alpha}$ as defined in Definition 4.2.*

1. *The (α, β)-exploration-replication policy behaves on Γ precisely as the $(0, \beta)$-exploration-replication policy does on $\Gamma^{+\alpha}$, i. e., the resulting migration rates ρ_{PQ} are identical.*

2. *If $\Phi^{+\alpha}(f) \leq (1+\epsilon)\,(\Phi^{+\alpha})^*$, then $\Phi(f) \leq (1+\epsilon)\,\Phi^* + \epsilon\,\alpha$.*

Proof. For the first statement, just observe that Equation (4.2) is equivalent to Equation (4.2) with $\alpha = 0$ and the latency functions ℓ_P and ℓ_Q substituted by $\ell_P + \alpha$ and $\ell_Q + \alpha$, respectively. For the second statement, observe that

$$\begin{aligned} \Phi(f) &= \Phi^{+\alpha}(f) - \alpha \\ &\leq (1+\epsilon)\,(\Phi^{+\alpha})^* - \alpha \\ &= (1+\epsilon)\,(\Phi^* + \alpha) - \alpha \\ &= (1+\epsilon)\,\Phi^* + \epsilon\,\alpha \ . \quad \square \end{aligned}$$

4.3 Convergence

Observe that the (α-β)-exploration-replication policy does *not* satisfy the smoothness condition required in Chapter 3. The migration probability $\mu(\ell_P, \ell_Q) = (\ell_P - \ell_Q)/\ell_P$ can exceed $\ell_P - \ell_Q$ by any factor if the denominator ℓ_P is sufficiently small. Thus, to prove convergence, we have to extend the proof of Lemma 3.3 to our dynamics. Again, the proof proceeds by showing that for consecutive flows f and f' the true potential gain $\Phi(f') - \Phi(f)$ is at least half of the virtual potential gain $\mathcal{V}(f, f')$. To that end, we give bounds on the error terms U_e by which the true potential gain differs from the virtual potential gain. The proof makes use of the fact that latency functions have elasticity d to bound the error terms caused by proportional sampling. The contribution to the error terms caused by uniform sampling can be bounded since β is chosen according to Equation (4.1). Recall the definition of V_{PQ}, U_e, and $\mathcal{V}$ from Equations (3.5), (3.6), and (3.7).

Lemma 4.4. *Consider an instance Γ, a flow vector $f \in \mathcal{F}_\Gamma$, and a flow vector $f' \in \mathcal{F}_\Gamma$ generated by the (α, β)-exploration-replication policy from f in one round. Then,*

$$\Delta\Phi = \Phi(f') - \Phi(f) \leq \frac{1}{2} \sum_{P,Q \in \mathcal{P}} \rho_{PQ}\,(\ell_Q - \ell_P) = \frac{1}{2} \cdot \mathcal{V}(f, f') \ .$$

Proof. By Lemma 3.2,

$$\Phi(f') - \Phi(f) = \sum_{P,Q \in \mathcal{P}} V_{PQ} + \sum_{e \in E} U_e \ , \tag{4.3}$$

where $V_{PQ} = \rho_{PQ}\,(\ell_Q - \ell_P)$ is the *virtual potential* gain of the agents migrating from P to Q and $U_e = \int_{f_e}^{f'_e} (\ell_e(u) - \ell_e(f_e))\,du$ is the *error term* caused by a latency increase or decrease of resource e. To prove the claim, we distribute the virtual potential gain terms V_{PQ} among all resources $e \in P \cup Q$. More precisely, for $e \in P$ and $e' \in Q$, we define

$$V_{PQ}^{e} = \rho_{PQ} \cdot \left(\left(\frac{1}{8}(\ell_Q - \ell_P) \cdot \frac{\ell_e}{\ell_P} \right) + \left(\frac{1}{8}(\ell_Q - \ell_P) \cdot \frac{1}{|P|} \right) \right) \tag{4.4}$$

and

$$V_{PQ}^{e'} = \rho_{PQ} \cdot \left(\left(\frac{1}{8}(\ell_Q - \ell_P) \cdot \frac{\ell_{e'}}{\ell_Q} \right) + \left(\frac{1}{8}(\ell_Q - \ell_P) \cdot \frac{1}{|Q|} \right) \right) \ .$$

Note that

$$\begin{aligned} \sum_{e \in P, e \in Q} V_{PQ}^{e} &= \frac{\rho_{PQ} \cdot (\ell_P - \ell_Q)}{8} \cdot \left(\sum_{e \in P} \left(\frac{\ell_e}{\ell_P} + \frac{1}{|P|} \right) + \sum_{e \in Q} \left(\frac{\ell_e}{\ell_Q} + \frac{1}{|Q|} \right) \right) \\ &= \frac{\rho_{PQ} \cdot (\ell_P - \ell_Q)}{8} \cdot \left(\frac{\ell_P}{\ell_P} + \frac{|P|}{|P|} + \frac{\ell_Q}{\ell_Q} + \frac{|Q|}{|Q|} \right) \\ &= \frac{V_{PQ}}{2} \ , \end{aligned}$$

i. e., these terms consume precisely half of the virtual potential gain. Substituting this into Equation (4.3),

$$\Delta\Phi = \frac{1}{2} \sum_{P,Q \in \mathcal{P}} V_{PQ} + \sum_{e \in E} \left(U_e + \sum_{P,Q \ni e} V_{PQ}^{e} \right) \ . \tag{4.5}$$

We now compare the terms V_{PQ}^{e} to the error terms U_e. We show that each term of the second sum is non-positive which concludes the proof of the lemma.

Consider an arbitrary resource e with $f'_e > f_e$. We consider the pairs of paths (P_j, Q_j) with $e \in Q_j$ in ascending order of $(\ell_{P_j} - \ell_{Q_j})/(\ell_{P_j} + \alpha)$ with $j \in \{1, \ldots, p\}$ where $p \le |\mathcal{P}|^2$ it the number of pairs of paths. Let

$$f_e^0 = f_e \qquad \text{and} \qquad f_e^j = f_e^{j-1} + \rho_{P_j Q_j} \quad \text{for } j > 0 \ ,$$

i. e., f_e^j is the flow on edge e counting the original flow and the contribution of the first j migrations only. By induction on n we show that

$$\int_{f_e}^{f_e^n} (\ell_e(u) - \ell_e(f_e))\,du \le -\sum_{j=1}^{n} V_{P_j Q_j}^{e} \ ,$$

which, for $n = p$, implies $U_e + \sum_{PQ} V_{PQ}^e \leq 0$ and hence the second sum in Equation (4.5) is non-positive concluding the proof of the lemma. The statement certainly holds for $n = 0$, and, by induction hypothesis,

$$\int_{f_e}^{f_e^n} (\ell_e(u) - \ell_e(f_e))\, du \quad \leq \quad -\sum_{j=1}^{n-1} V_{P_j Q_j}^e + \int_{f_e^{n-1}}^{f_e^n} (\ell_e(u) - \ell_e(f_e))\, du \ .$$

Hence, it remains to show that

$$\int_{f_e^{n-1}}^{f_e^n} (\ell_e(u) - \ell_e(f_e))\, du \leq -V_{P_n Q_n}^e \ . \tag{4.6}$$

There are two cases.

Case 1. $f_e \geq \beta$, i. e., the migration towards e is dominated by proportional sampling. Denote the flow on edge e belonging to commodity i by $f_{e,i}$ and consider the total volume of agents sampling a path containing edge e. This volume is bounded from above by

$$\sum_{i \in [k]} r_i \cdot \left(\beta + (1 - \beta) \cdot \frac{f_{e,i}}{r_i} \right) = \beta + (1 - \beta) \cdot f_e \leq 2\, f_e \ .$$

Multiplying this with the activation and migration probability, which is, for the first n path pairs, bounded by $\lambda / d \cdot (\ell_{P_n} - \ell_{Q_n}) / (\ell_{P_n} + \alpha)$, we obtain

$$f_e^n - f_e \leq \frac{2\, f_e}{32\, d} \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{\ell_{P_n}(f) + \alpha} \ . \tag{4.7}$$

We use that ℓ_e has elasticity d and apply Fact 4.1 to obtain

$$\ell_e(f_e^n) - \ell_e(f_e) \leq 2\, d \cdot \frac{f_e^n - f_e}{f_e} \cdot \ell_e(f_e) \ .$$

Substituting Equation (4.7) into this and using $\alpha \geq 0$ yields

$$\ell_e(f_e^n) - \ell_e(f_e) \leq d \cdot \frac{\ell_e(f_e)}{f_e} \cdot \frac{1}{8\, d} \cdot f_e \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{\ell_{P_n}(f)}$$

and hence

$$\begin{aligned} \rho_{P_n Q_n} \cdot (\ell_e(f_e^n) - \ell_e(f_e)) \quad & \leq \quad \rho_{P_n Q_n} \cdot \frac{1}{8} \cdot \frac{\ell_e(f_e)}{\ell_{P_n}(f)} \cdot (\ell_{P_n}(f) - \ell_{Q_n}(f)) \\ & \leq \quad -V_{P_n Q_n}^e \end{aligned}$$

which proves Equation (4.6).

Case 2. $f_e < \beta$, i. e., the migration towards e is dominated by uniform sampling. This time, the volume of agents sampling a path containing e is at most

$$\beta + (1-\beta) \cdot f_e \leq 2\,\beta \ .$$

Multiplying this with the migration probability, the increase of flow due to the first n paths is at most

$$f_e^n - f_e \leq \beta \cdot \frac{2}{32\,d} \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{\ell_{P_n}(f) + \alpha} \ .$$

By choice of β according to Equation (4.1), the maximum derivative of ℓ_e in the range $[0, \beta]$ is at most $(\ell_{P_n}(0) + \alpha)/(D\,\beta)$ and since $f_e^n \leq \beta + 2\beta/(32\,d) = \beta(1 + 1/(16\,d))$, the derivative in the range $[0, f_e^n]$ is at most twice as large. Altogether,

$$\begin{aligned}
\ell_e(f_e^n) - \ell_e(f_e) &\leq (f_e^n - f_e) \cdot \frac{2\,(\ell_{P_n}(0) + \alpha)}{D\,\beta} \\
&\leq \frac{1}{8\,d} \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{|P_n| \cdot (\ell_{P_n}(f) + \alpha)} \cdot (\ell_{P_n}(0) + \alpha) \\
&\leq \frac{1}{8} \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{|P_n|} \ .
\end{aligned}$$

Therefore,

$$\begin{aligned}
\rho_{P_n Q_n} \cdot (\ell_e(f_e^n) - \ell_e(f_e)) &\leq \rho_{P_n Q_n} \cdot \frac{1}{8} \cdot \frac{\ell_{P_n}(f) - \ell_{Q_n}(f)}{|P_n|} \\
&\leq -V^e_{P_n Q_n}
\end{aligned}$$

which, again, proves Equation (4.6).

As in the proof of Lemma 3.3, resources with $f'_e < f_e$ can be treated similarly with the roles of P and Q exchanged. Here, we need to consider only Case 1 since the rate at which agents leave a path is always proportional to its current flow. Consequently, Equation (4.7) holds (with reversed inequality) irrespective of whether $f_e \leq \beta$ or $f_e > \beta$. □

From this lemma, convergence follows as in Corollary 3.4.

Corollary 4.5. *For any initial population $f_0 \in \mathcal{F}_\Gamma$ and $\beta > 0$ satisfying Equation (4.1), the (α, β)-exploration-replication policy converges towards the set of Wardrop equilibria.*

4.4 Convergence Time for Single-Commodity Instances

In this section, we consider the case of symmetric instances, i. e., the number of commodities is $k = 1$. We will first derive upper bounds for the time of convergence towards approximate equilibria in a bicriterial sense and proceed by giving upper bounds for the number of rounds until the potential is close to the optimum potential.

4.4.1 Bicriteria Approximation

In this section we will use the following bicriterial definition of approximate equilibria.

Definition 4.4 ((δ,ϵ)-equilibrium). *For a symmetric instance Γ and a flow vector $f \in \mathcal{F}_\Gamma$ let $\mathcal{P}^+(\delta) = \{P \in \mathcal{P} \mid \ell_P(f) \geq (1+\delta)\, L(f)\}$ denote the set of δ-*expensive paths *and let $\mathcal{P}^-(\delta) = \{P \in \mathcal{P} \mid \ell_P(f) \leq (1-\delta)\, L(f)\}$ denote the set of δ-*cheap paths*. The population f is at a (δ,ϵ)-equilibrium iff*

$$\sum_{P \in \mathcal{P}^+(\delta) \cup \mathcal{P}^-(\delta)} f_P \leq \epsilon \ .$$

We write $\mathcal{P}^+$ and $\mathcal{P}^-$ if δ is clear from the context. Approximate equilibria in this sense are again transient for the (α,β)-exploration-replication policy. Hence, we bound the total number of rounds that are not at an approximate equilibrium.

Theorem 4.6. *Consider a symmetric instance Γ, an initial flow vector $f_0 \in \mathcal{F}_\Gamma$ and constants $\alpha, \beta \geq 0$. For the (α, β)-exploration-replication policy, the number of rounds in which the population is not at a (δ,ϵ)-equilibrium w. r. t. $\Gamma^{+\alpha}$ (Definition 4.2) is bounded from above by*

$$\mathcal{O}\left(\frac{d}{\epsilon\, \delta^2} \ln\left(\frac{\Phi(f_0) + \alpha}{\Phi^* + \alpha}\right)\right) \ .$$

In particular, this bound holds for $\alpha = \beta = 0$ (and hence $\Gamma^{+\alpha} = \Gamma$).

Proof. It is sufficient to consider the case $\alpha = 0$ (and $\Phi^* > 0$). To see this, assume that the theorem is valid for $\alpha = 0$ and consider an instance $\hat{\Gamma}$ and some $\hat{\alpha} > 0$. By Fact 4.3, applying the lemma to $\Gamma = \hat{\Gamma}^{+\hat{\alpha}}$ and $\alpha = 0$ yields the assertion of the lemma applied to $\hat{\Gamma}$ and $\hat{\alpha}$.

We estimate the virtual potential gain of a round that starts with a population that is not at a (δ,ϵ)-equilibrium. Recall that the virtual potential gain is the difference between the potential of two consecutive rounds assuming that the latency functions were fixed at the beginning of the first

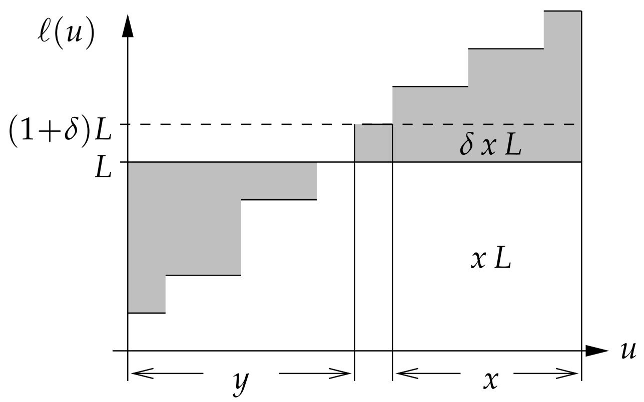

Figure 4.2: The figure depicts the distribution of the agents' latencies. The shaded areas are of the same size. Since the left area represents the conditional expectation of the difference between L and the latency of an agent with latency in the range $[0, L]$, this expectation has a value of at least $\delta\, x\, L$.

round. By Lemma 4.4, the true potential gain is at least half of the virtual potential gain. As long as we are not at a (δ,ϵ)-equilibrium, at least one of the following cases holds.

Case 1. At least a volume of $\epsilon/2$ agents utilises δ-expensive paths.

Case 1a. At least a volume of $x \geq 1/2$ agents utilises δ-expensive paths. Let $y > 0$ denote the volume of agents utilising paths with latency at most L. Consider a random agent that is sampled by proportional sampling and let Y denote the random variable that represents its current latency. Let $\Delta = \mathbb{E}\left[L - Y \mid Y \leq L\right]$ denote the conditional expectation of the difference between Y and L. For an illustration see Figure 4.2. We have

$$\Delta = L - \frac{1}{y}\left(\sum_{P:\ell_P \leq L} f_P\, \ell_P\right) \quad . \tag{4.8}$$

Furthermore, by definition of x,

$$\sum_{P:\ell_P > (1+\delta)L} f_P\, \ell_P \geq \delta\, x\, L + x\, L \quad . \tag{4.9}$$

Finally, clearly

$$\sum_{P:L < \ell_P \leq (1+\delta)L} f_P\, \ell_P \geq (1 - x - y) \cdot L \quad . \tag{4.10}$$

Substituting Equations (4.8), (4.9), and (4.10) into the definition of L,

$$\begin{aligned} L &= \sum_{P:\ell_P \le L} f_P\,\ell_P + \sum_{P:L<\ell_P\le(1+\delta)L} f_P\,\ell_P + \sum_{P:\ell_P>(1+\delta)L} f_P\,\ell_P \\ &\ge y\cdot(L-\Delta) + (1-x-y)\cdot L + \delta\,x\,L + x\,L \\ &= L - y\,\Delta + \delta\,x\,L \end{aligned}$$

implying $\Delta \ge (\delta\,x/y)\,L$. All agents with latency at least $(1+\delta)\,L$ that sample a path with latency at most L migrate to the sampled path with probability at least $\lambda\cdot\delta/(d\,(1+\delta))$. The probability to sample such a path is at least $(1-\beta)\,y$. Their expected latency gain which is equivalent to their infinitesimal contribution to the virtual potential gain is at least Δ. In total there is a volume of $x \ge 1/2$ such agents. Altogether, the virtual potential gain is

$$\mathcal{V} \le -x\cdot\frac{\lambda\,(1-\beta)\,y\,\delta}{d\,(1+\delta)}\cdot\Delta \le -\frac{\lambda\,(1-\beta)\,x^2\,\delta^2}{d\,(1+\delta)}L \le -\frac{\lambda\,\delta^2}{8\,d}L\ .$$

Case 1b. At least a volume of $\epsilon/2$ but at most a volume of $1/2$ agents utilises δ-expensive paths. Then, the virtual potential gain of the agents leaving δ-expensive paths is

$$\begin{aligned} \mathcal{V} &\le -\sum_{P\in\mathcal{P}^+} f_P \sum_{Q\notin\mathcal{P}^+} \rho_{PQ}(\ell_P-\ell_Q) \\ &\le -\frac{\lambda}{d}\sum_{P\in\mathcal{P}^+} f_P \sum_{Q\notin\mathcal{P}^+}\left((1-\beta)f_Q+\frac{\beta}{|\mathcal{P}|}\right)\frac{(\ell_P-\ell_Q)^2}{\ell_P}\ . \end{aligned}$$

Omitting the term $\beta/|\mathcal{P}|$, substituting $\ell_P \ge (1+\delta)\,L$ and $(1-\beta)/(1+\delta) \ge 1/2$, and applying Jensen's inequality to the last sum yields

$$\begin{aligned} \mathcal{V} &\le -\frac{\lambda}{2\,d\,L}\sum_{P\in\mathcal{P}^+} f_P\left(\sum_{Q\notin\mathcal{P}^+} f_Q(L+\delta L-\ell_Q)\right)^2\cdot\frac{1}{\sum_{Q\notin\mathcal{P}^+} f_Q} \\ &\le -\frac{\lambda}{2\,d\,L}\sum_{P\in\mathcal{P}^+} f_P\left(L\sum_{Q\notin\mathcal{P}^+} f_Q + \delta\,L\sum_{Q\notin\mathcal{P}^+} f_Q - \sum_{Q\notin\mathcal{P}^+} f_Q\,\ell_Q\right)^2\ . \end{aligned}$$

Note that $\sum_{Q\notin\mathcal{P}^+} f_Q\ell_Q \le L\sum_{Q\notin\mathcal{P}^+} f_Q$ since the sum omits the terms of the expensive paths $Q \in \mathcal{P}^+$. Hence,

$$\begin{aligned}
\mathcal{V} &\leq -\frac{\lambda}{2\,d\,L} \sum_{P\in\mathcal{P}^+} f_P \left(\delta\, L \sum_{Q\notin\mathcal{P}^+} f_Q\right)^2 \\
&\leq -\frac{\lambda\,\delta^2\,L^2}{2\,d\,L} \sum_{P\in\mathcal{P}^+} f_P \left(\sum_{Q\notin\mathcal{P}^+} f_Q\right)^2 \\
&\leq -\frac{\lambda\,(\epsilon/2)(1/2)^2\,\delta^2}{2\,d} L \\
&= -\frac{\lambda\,\epsilon\,\delta^2}{16\,d} L \ .
\end{aligned}$$

In the last inequality we used our assumptions that $\sum_{Q\notin\mathcal{P}^+} f_Q \geq 1/2$ and $\sum_{P\in\mathcal{P}^+} \geq \epsilon/2$.

Case 2. At least a volume of $\epsilon/2$ agents utilises δ-cheap paths.

Case 2a. At least a volume of $x \geq 1/2$ agents utilises δ-cheap paths. Let $y > 0$ denote the volume of agents utilising paths with latency at least L. Consider a random agent that is activated in this round and let Y denote the random variable that represents its current latency. Let $\Delta = \mathbb{E}\left[Y - L \mid Y \geq L\right]$ denote the conditional expectation of the difference by which Y exceeds L. Using similar arguments as in case 1a, we obtain

$$\begin{aligned}
\Delta &= \frac{1}{y}\left(\sum_{P:\ell_P\geq L} f_P\,\ell_P\right) - L \ , \\
\sum_{P:\ell_P<(1-\delta)L} f_P\,\ell_P &\leq x\,L - \delta\,x\,L \ \text{, and} \\
\sum_{P:(1-\delta)\leq\ell_P<L} f_P\,\ell_P &\leq (1-x-y)\cdot L \ .
\end{aligned}$$

This yields

$$\begin{aligned}
L &\leq y\cdot(\Delta + L) + (1-x-y)\cdot L + x\,L - \delta\,x\,L \\
&= L + \Delta\,y - \delta\,x\,L \ ,
\end{aligned}$$

again implying $\Delta \geq (\delta\,x/y)\cdot L$. An agent with latency at least L that samples a path with latency at most $(1-\delta)\,L$ migrates to this path with probability at least $\lambda\,\delta/d$. The probability to sample such a path is at least $(1-\beta)\,x$, and there is a volume of y such agents. Their expected latency gain is Δ. Hence, their virtual potential gain is

$$\mathcal{V} \leq -y \cdot \frac{\lambda\,(1-\beta)\,x\,\delta}{d} \cdot \Delta \leq -\frac{\lambda\,\delta^2}{8\,d} L \ .$$

Case 2b. At least $\epsilon/2$ but at most $1/2$ agents utilise δ-cheap paths. We can treat the virtual potential gain of agents moving towards δ-cheap paths in a way similar to case 1b.

$$\begin{aligned}\mathcal{V} &\leq -\sum_{Q\notin\mathcal{P}^-} f_Q \sum_{P\in\mathcal{P}^-} \rho_{QP}(\ell_Q - \ell_P)\\ &\leq -\frac{\lambda}{d}\sum_{P\in\mathcal{P}^-}\left((1-\beta)f_P + \frac{\beta}{|\mathcal{P}|}\right)\sum_{Q\notin\mathcal{P}^-} f_Q\,\ell_Q\left(\frac{\ell_Q-\ell_P}{\ell_Q}\right)^2 .\end{aligned}$$

Now, we use that $\ell_P \leq L - \delta\,L$ and apply Jensen's inequality:

$$\begin{aligned}\mathcal{V} &\leq -\frac{\lambda}{2\,d}\sum_{P\in\mathcal{P}^-} f_P\left(\sum_{Q\notin\mathcal{P}^-} f_Q\,\ell_Q\frac{\ell_Q-\ell_P}{\ell_Q}\right)^2\frac{1}{\sum_{Q\notin\mathcal{P}^-} f_Q\,\ell_Q}\\ &\leq -\frac{\lambda}{2\,d\,L}\sum_{P\in\mathcal{P}^-} f_P\left(\sum_{Q\notin\mathcal{P}^-} f_Q(\ell_Q - L + \delta\,L)\right)^2\\ &\leq -\frac{\lambda}{2\,d\,L}\sum_{P\in\mathcal{P}^-} f_P\left(\sum_{Q\notin\mathcal{P}^-} f_Q\ell_Q - L\sum_{Q\notin\mathcal{P}^-} f_Q + \delta\,L\sum_{Q\notin\mathcal{P}^-} f_Q\right)^2\end{aligned}$$

Note that $\sum_{Q\notin\mathcal{P}^-} f_Q\ell_Q \geq L\sum_{Q\notin\mathcal{P}^-} f_Q$ since the sum omits the terms of the paths $Q\in\mathcal{P}^-$ with small latency. Hence,

$$\begin{aligned}\mathcal{V} &\leq -\frac{\lambda}{2\,d\,L}\sum_{P\in\mathcal{P}^-} f_P\left(\delta\,L\cdot\sum_{Q\notin\mathcal{P}^-} f_Q\right)^2\\ &\leq -\frac{\lambda\,\delta^2\,L^2}{2\,d\,L}\sum_{P\in\mathcal{P}^-} f_P\left(\sum_{Q\notin\mathcal{P}^-} f_Q\right)^2\\ &\leq -\frac{\lambda\,(\epsilon/2)(1/2)^2\,\delta^2}{2\,d}L\\ &\leq -\frac{\lambda\,\epsilon\,\delta^2}{16\,d}L\ .\end{aligned}$$

In all four cases we have $\mathcal{V} \leq -\frac{\lambda\,\epsilon\,\delta^2 L}{16\,d}$. Due to Lemma 4.4 and Fact 4.2, the true potential gain is

$$\Delta\Phi \leq \frac{\mathcal{V}}{2} \leq -\frac{\lambda\,\epsilon\,\delta^2}{32\,d}\Phi.$$

Let $\Phi(t)$ denote the potential in the t-th round that is not at a (δ,ϵ)-equilibrium. Then,

$$\Phi(t) \leq \Phi(f_0)\cdot\left(1-\frac{\lambda\,\epsilon\,\delta^2}{32\,d}\right)^t .$$

Since Φ is lower bounded by Φ^* we obtain the desired upper bound on the number of unbalanced phases. □

4.4.2 Approximation of the Potential

Since (δ,ϵ)-equilibria are transient, we propose an alternative definition of approximate equilibria in this section. We derive bounds on the time to reach a flow which approximates the optimal potential Φ^* up to a factor of at most $(1+\epsilon)$. Clearly, approximate equilibria in this sense are persistent. First note that if $\Phi^* = 0$, the potential cannot be approximated up to a relative factor unless exactly optimised. We therefore allow an additional deviation by an additive term $\epsilon\,\alpha$, i. e., we want to reach a population f with potential at most $\Phi(f) \leq (1+\epsilon)\,\Phi^* + \epsilon\,\alpha$. We start with several lemmas that can be applied to symmetric games in general and proceed by analysing the symmetric singleton case in which $|P| = D = 1$ for all $P \in \mathcal{P}$ and the general symmetric case separately.

The first lemma shows that as long as we are at a (δ,ϵ)-equilibrium, the volume of agents moving within one round is small.

Lemma 4.7. *For a flow $f \in \mathcal{F}_\Gamma$ at (δ,ϵ)-equilibrium, the total volume of flow moved by the (α,β)-exploration-replication policy within one round is at most*

$$\sum_{P,Q\in\mathcal{P}} \rho_{PQ}(f) \leq \frac{\lambda}{d}\cdot(2\,\epsilon + 2\,\delta + \beta) \enspace .$$

Proof. We can write the above sum as

$$\begin{aligned}\sum_{P,Q\in\mathcal{P}} \rho_{PQ} &= \sum_{P\in\mathcal{P}^+}\sum_{Q\in\mathcal{P}} \rho_{PQ} + \sum_{P\notin\mathcal{P}^+}\sum_{Q\notin\mathcal{P}^-} \rho_{PQ} + \sum_{P\notin\mathcal{P}^+}\sum_{Q\in\mathcal{P}^-} \rho_{PQ} \\ &\leq \frac{\lambda}{d}\,(\epsilon + 2\,\delta + (\epsilon+\beta)) \enspace .\end{aligned}$$

The first estimate holds since at most a volume of ϵ agents utilises paths in $\mathcal{P}^+$, the second holds since the probability to migrate from a path with latency at most $(1+\delta)\,L$ to a path with latency at least $(1-\delta)\,L$ is at most $(1+\delta-(1-\delta))/(1+\delta) \leq 2\,\delta$, and the third holds since the probability to sample a path in $\mathcal{P}^-$ is at most $\beta + (1-\beta)\,\epsilon$. □

Next we show that the value of the average latency $L = \sum_{P\in\mathcal{P}} \ell_p\, f_P = \sum_{e\in E} \ell_e\, f_e$ does not decrease much within one round unless the potential also decreases significantly. To that end we show that the drift of the average latency and the drift of the potential are coupled. Every resource $e \in E$ contributes to both the potential Φ and the average L. The contribution to Φ is the area under the graph of the latency function between 0 and f_e, and the contribution to the average is the area of the rectangle between the origin and the point $(f_e, \ell_e(f_e))$. When agents are removed from a resource, the contribution of the resource to the latency average is reduced by an

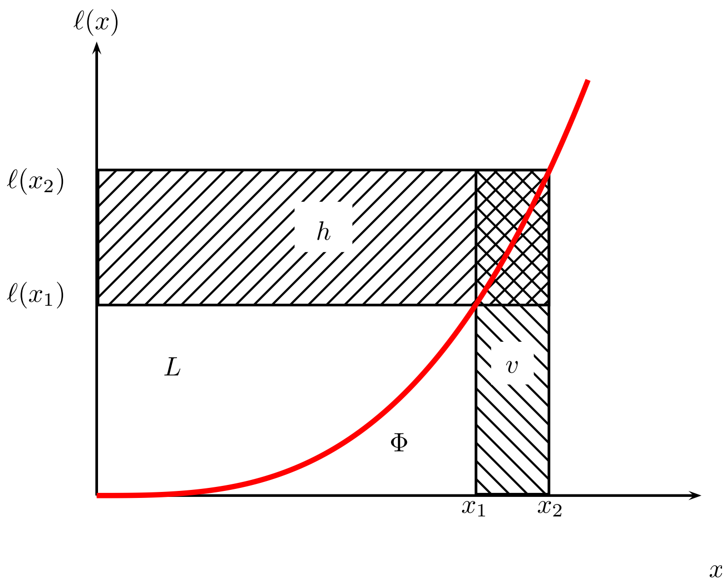

Figure 4.3: This figure illustrates the contribution to the average L and the potential Φ when the flow on a link changes from x_1 to x_2.

L-shaped area consisting of a horizontal and a (partially overlapping) vertical rectangle, denoted h and v, respectively (see Figure 4.3). The following lemma states that the horizontal bar is at most d times greater than the vertical bar.

Lemma 4.8. *Consider a function $\ell(\cdot)$ with elasticity d and $x_1, x_2 \in [0, 1]$ with $x_1 < x_2$. Let $v = (x_2 - x_1)\,\ell(x_2)$ and $h = (\ell(x_2) - \ell(x_1))\,x_2$. Then, $h \le d\,v$.*

Proof. Since ℓ has elasticity d, we know that the slope of ℓ at x_2 is at most $m \le (\ell(x_2)/x_2)\,d$. Consequently, $\ell(x_1) \ge \ell(x_2) - m\,(x_2 - x_1)$. Substituting both expressions into the definition of v shows that

$$\begin{aligned} h &= (\ell(x_2) - \ell(x_1))\,x_2 \\ &\le (\ell(x_2) - \ell(x_2) + m\,(x_2 - x_1)) \cdot x_2 \\ &\le \frac{\ell(x_2)}{x_2} \cdot d \cdot (x_2 - x_1) \cdot x_2 \\ &= d \cdot v\ . \quad \Box \end{aligned}$$

Lemma 4.9. *Consider a symmetric instance Γ, a flow $f \in \Gamma$ at (δ,ϵ)-equilibrium, and a flow f' generated from f by the (α,β)-exploration-replication policy within one round. If $L(f') \le L(f) - \Delta$ with $\Delta > 10\,\lambda \cdot (2\,\epsilon + 2\,\delta + \beta)\,L$, then $\Phi(f') \le \Phi(f) - \Delta/(10\,(d+1))$.*

Proof. First consider the contribution of agents leaving paths with latency at most $3\,L$. Due to Lemma 4.7, their total volume is at most $(\lambda/d)\,(2\,\epsilon + 2\,\delta + \beta)$. Hence, the contribution to the potential decrease (the "vertical bars") caused by removing these agents and pessimistically ignoring the fact that they are re-added to other paths sum up to at most $3\,(\lambda/d) \cdot (2\,\epsilon + 2\,\delta + \beta)\,L$. Due to Lemma 4.8, the total decrease of the average latency caused by these agents is at most $\Delta_1 \le (d+1) \cdot 3\,\lambda/d \cdot (2\,\epsilon + \beta + 2\,\delta)L \le 6\,\lambda\,(2\,\epsilon + \beta + 2\,\delta)\,L$.

Now, consider the agents leaving paths with latency above $3\,L$. These add an additional contribution of Δ_2 to the decrease of the average. Since f is at a (δ,ϵ)-equilibrium, the probability to sample a path P with $\ell_P \le (1+\delta)\,L$ is at least $(1-\beta) \cdot (1-\epsilon)$. The migration probability for such a path is greater than the migration probability for a path with higher latency, independent of the agent's origin path. Therefore, the conditional probability that a migrating agent coming from a path P with $\ell_P > 3\,L$ migrates to a path Q with $\ell_Q \le (1+\delta)\,L$ is at least $(1-\epsilon-\beta)$. Hence, an agent coming from such a path P has an expected latency gain of at least $(1-\epsilon-\beta)\,(\ell_P - (1+\delta)\,L) \ge \ell_P/2$. This latency gain equals the agent's infinitesimal contribution to the virtual potential gain. Again, by Lemma 4.8, its infinitesimal contribution to the decrease of the average latency is at most $\ell_P \cdot (d+1)$. Therefore, the contribution Δ_2 to the decrease of the average is separated from the virtual potential gain by no more than a factor of $2 \cdot (d+1)$ and consequently the virtual potential gain of these agents is at least $\Delta_2/(2\,(d+1))$. By Lemma 4.4, the true potential gain is at least $\Delta_2/(4\,(d+1))$.

Now, consider the total decrease of the average latency $\Delta = \Delta_1 + \Delta_2$. There are two cases.

Case 1. $\Delta_2 \le 4\,\lambda(2\,\epsilon + 2\,\delta + \beta)\,L$. Then, $\Delta = \Delta_1 + \Delta_2 \le 10\,\lambda(2\,\epsilon + 2\,\delta + \beta)\,L$, our claim.

Case 2. $\Delta_2 > 4\,\lambda(2\,\epsilon + 2\,\delta + \beta)\,L$. Then, since $\Delta_2 \ge \Delta \cdot 4/10$,

$$\Phi(f') - \Phi(f) \le -\frac{\Delta_2}{4\,(d+1)} \le -\frac{\Delta}{10\,(d+1)} \ .$$

Again, this is the assertion of the lemma. □

Our analysis of the time of convergence proceeds by constructing a flow vector in which the latencies of the cheapest path and the most expensive used path are close to the average latency and consequently close to each other. The following lemma shows that such a configuration in which all paths deviate by no more than a fixed fraction of L from the average latency of their commodity, are also approximations of the optimal potential.

Definition 4.5 (δ-bounded)**.** *A flow vector $f \in \mathcal{F}_\Gamma$ is* δ-bounded *if for every commodity $i \in [k]$ and for every $P \in \mathcal{P}_i$ it holds that $\ell_P(f) \geq L_i - \delta\, L$ and, in addition, if $f_P > 0$, $\ell_P(f) \leq L_i + \delta\, L$.*

Lemma 4.10. *For every instance and δ-bounded flow vector $f \in \mathcal{F}_\Gamma$,*

$$\frac{\Phi(f)}{\Phi^*} \leq 1 + \mathcal{O}\left(\delta\, d\right) \ .$$

Proof. Consider the δ-bounded flow f for the instance Γ. For every $i \in [k]$ let $\ell^i_{\max}(f) = \max_{P \in \mathcal{P}_i, f_P > 0} \ell_P(f)$. We extend every strategy $P \in \mathcal{P}_i$ with $\ell_P(f) \leq \ell^i_{\max}$ by a new resource with constant latency $\delta_P = \ell^i_{\max}(f) - \ell_P(f)$. Since f is δ-bounded, we can be sure that $\delta_P \leq 2\,\delta\, L$. Let Γ' denote the thus obtained instance. Since in Γ' the latencies of all used paths are minimal, f is at a Wardrop equilibrium for Γ' with $\Phi_{\Gamma'}(f) \geq \Phi_\Gamma(f)$.

Now, consider the minimal potential in Γ, denoted by Φ^*_Γ. We have to show that $\Phi_{\Gamma'}(f) \leq \Phi^*_\Gamma + 2\,\delta\, L$ since this and Fact 4.2 imply that $\Phi_\Gamma(f) \leq \Phi^*_\Gamma + 2\,\delta\,(d+1)\,\Phi_\Gamma(f)$ which is the assertion of our theorem. To see the claim, note that for every $\tilde{f}$ it holds that $\Phi_{\Gamma'}(\tilde{f}) \leq \Phi_\Gamma(\tilde{f}) + 2\,\delta\, L$. Also, the constraints for $\tilde{f}$ under which $\Phi_{\Gamma'}$ and Φ_Γ are to be minimised are identical, i. e., every $\tilde{f}$ feasible for Γ is also feasible for Γ' and vice versa. By optimality of f in Γ', we have $\Phi_{\Gamma'}(f) \leq \Phi_{\Gamma'}(\tilde{f}) \leq \Phi_\Gamma(\tilde{f}) + 2\,\delta\, L$ for any $\tilde{f}$ and specifically for an $\tilde{f}$ that satisfies $\Phi_\Gamma(\tilde{f}) = \Phi^*_\Gamma$. ☐

Let us remark that our policy will not reach a δ-bounded state, since the population on high-latency paths will not vanish completely. However, our policy will reach a state at which the population on such paths is so small that removing them will result in a δ-bounded flow without changing the potential significantly.

Finally, the following lemma guarantees that the optimal potential does not increase significantly if the total flow demand of an instance is increased slightly.

Lemma 4.11. *For any $x \in [0, 1/(4\,d)]$, let $\Gamma = (G, \{(s,t,1)\}, (\ell)_{e \in E})$ and $\tilde{\Gamma} = (G, \{(s,t,1+x)\}, (\ell)_{e \in E})$ be symmetric instances, and let Φ^* and $\tilde{\Phi}^*$ denote their respective minimal potential values. Then, $\tilde{\Phi}^* \leq (1 + 3\,d\,x)\Phi^*$.*

Proof. Let $f \in \mathcal{F}_\Gamma$ denote a flow with $\Phi(f) = \Phi^*$ and let $\tilde{f} = (1+x) \cdot f$. Note that $\tilde{f}$ is feasible for $\tilde{\Gamma}$. Then we have $\tilde{f}_e = (1+x) f_e$ and $\ell_e(\tilde{f}) \leq$

$(1 + 2\,x\,d) \cdot \ell_e(f_e)$. Thus,

$$\begin{aligned} \Phi(\tilde{f}) - \Phi(f) &\leq \sum_{e \in E} \int_{f_e}^{(1+x)f_e} \ell_e(u) du \\ &\leq \sum_{e \in E} (1 + 2\,x\,d)\, \ell_e(f_e)\, f_e\, x \\ &\leq (1 + 2\,x\,d)\, x\, L(f) \\ &\leq 3\,d\,x\,\Phi(f) \end{aligned}$$

by Fact 4.2. Hence, $\tilde{\Phi}^* \leq \Phi(\tilde{f}) \leq (1 + 3\,d\,x)\,\Phi^*$. □

4.4.2.1 Symmetric Singleton Games

For the case in which every path utilises only one resource, i. e., for all $P \in \mathcal{P}$, $|P| = D = 1$, we can now show convergence towards potential approximations.

Lemma 4.12. *Consider a symmetric singleton instance Γ and the $(0, \beta)$-exploration-replication policy for $0 < \beta \leq \epsilon^2/(d^3 \cdot \ln(|\mathcal{P}|\,d/\epsilon))$. For every $\epsilon > 0$ define the following constants:*

$$\delta = \frac{c\,\epsilon}{d}\,, \quad \delta' = \epsilon' = c' \frac{\delta^2}{d \ln(\delta\,|\mathcal{P}|/\beta)}\,, \text{and} \quad T = \frac{5\,d}{\lambda\,\delta} \ln\left(\frac{3\,\epsilon'\,d\,|\mathcal{P}|}{\delta\,\beta}\right),$$

where c and c' are positive constants independent of ϵ, β, and d to be defined in the proof. Then, in every phase consisting of T rounds starting with a flow vector f^0 at least one of the following events occurs:

(1) At least once in the phase the flow is not at a $(\delta'$,$\epsilon')$-equilibrium.

(2) At the end of the phase, $\Phi(f) \leq \Phi(f^0) - \Omega\left(\frac{\delta'}{d}\right) L(f^0)$.

(3) $\Phi(f^0) \leq (1 + \epsilon)\Phi^*$.

Proof. Consider a phase consisting of T rounds. Let us write $L_0 = L(f^0)$. We show that if events (1) and (2) do not occur, then event (3) must occur. Given that event (1) does not occur, Lemma 4.9 implies that the average $L(f(t))$ may decrease by at most a factor of $(1 - \mathcal{O}\left(\delta' + \epsilon' + \beta\right)) = (1 - \mathcal{O}\left(\delta'\right))$ per round or the potential decreases by at least $\Omega\left(L \cdot \delta'/d\right)$ which means that event (2) occurs. Consequently, if neither event (1) nor event (2) occur, the average L is at least

$$\begin{aligned} \left(1 - \mathcal{O}\left(\delta'\right)\right)^T \cdot L_0 &= \left(1 - \mathcal{O}\left(T \cdot \delta'\right)\right) \cdot L_0 \\ &= \left(1 - \mathcal{O}\left(\frac{c' \cdot \delta^2 \cdot d \cdot \ln(\epsilon'\,d\,|\mathcal{P}|/(\delta\,\beta))}{d \cdot \delta \cdot \ln(\delta\,|\mathcal{P}|/\beta)}\right)\right) \cdot L_0 \\ &= \left(1 - \mathcal{O}\left(c' \cdot \delta\right)\right) \cdot L_0\ . \end{aligned}$$

throughout the phase. In particular, $L(f(t)) \geq (1-\delta/4)\,L_0$ throughout the phase if c' is chosen small enough. Furthermore, at any time at least $1-\epsilon'$ agents utilise paths with latency at least $(1-\delta')(1-\delta/4)\,L_0 \geq (1-\delta/2)\,L_0$.

Consider a path P with latency below $(1-\delta)\,L_0$ and flow f_P. By our considerations above, at least $\lambda\, f_P\,\delta/(3\,d)$ agents migrate to this path in one round due to proportional sampling whereas at most $\lambda\, f_P\,(\epsilon'+\beta)/d$ agents may leave this path since the probability to sample a path with smaller latency is at most $\epsilon'+\beta$. Hence, the flow on such a path grows by at least

$$f_P\,\lambda\frac{\delta}{3\,d} - f_P\,\lambda\frac{\epsilon'+\beta}{d} \geq f_P\,\lambda\frac{\delta}{4\,d}$$

for small values of δ since $\epsilon', \beta = o(\delta)$. Hence, the population on every such path grows by a factor of at least $1+\lambda\,\delta/(4\,d)$ per round. Moreover, in the first round, the population grows to at least $\lambda\,\delta\,\beta/(3\,d\,|\mathcal{P}|)$ due to uniform sampling. Hence, provided that the latency of P stays below $(1-\delta)\,L_0$, at the end of the phase the population on such a path is at least

$$\begin{aligned}
\frac{\lambda\,\delta\,\beta}{3\,d\,|\mathcal{P}|}\left(1+\frac{\lambda\,\delta}{4\,d}\right)^{T} &= \frac{\lambda\,\delta\,\beta}{3\,d\,|\mathcal{P}|}\left(1+\frac{\lambda\,\delta}{4\,d}\right)^{\frac{5d}{\lambda\delta}\ln\left(\frac{3\epsilon'\,d\,|\mathcal{P}|}{\lambda\,\delta\,\beta}\right)} \\
&> \frac{\lambda\,\delta\,\beta}{3\,d\,|\mathcal{P}|}\exp\left(\ln\left(\frac{3\,\epsilon'\,d\,|\mathcal{P}|}{\lambda\,\delta\,\beta}\right)\right) \\
&= \epsilon' \;.
\end{aligned}$$

Here, we have used that $(1-1/X)^{X\cdot 5/4} > \exp(1)$ for $X \geq 2$ with $X = 4\,d/(\lambda\,\delta)$. This implies that at this point of time, either the population is no longer at an $(\delta'$,$\epsilon')$-equilibrium implying that event (1) occurs, or the latency of P has become at least $(1-\delta)L_0$. Hence, the latency of every path P with $\ell_P(f^0) < (1-\delta)L_0$ must grow to $(1-\delta)L_0$ at least once within the phase. Furthermore, by Lemma 4.7, the total amount of flow that is moved within the entire phase is at most $v = T\cdot\lambda\,(2\,\epsilon' + 2\,\delta' + \beta)/d$. Note that the arguments used so far do neither rely on the assumption that Γ is a singleton instance nor on the definition of δ.

Now, consider a flow vector f' obtained from f^0 by reducing the flow on any edge $e \in E$ with $\ell_e(f^0) > (1+\delta')L_0$ until either $\ell_e(f'_e) = (1+\delta')\,L_0$ or $f'_e = 0$. Then,

$$\begin{aligned}
\Phi(f^0) - \Phi(f') &\leq \sum_{e:\ell_e \geq (1+\delta')L_0} \int_0^{f_e^0} \ell_e(u)\,du \\
&\leq \sum_{e:\ell_e \geq (1+\delta')L_0} \ell_e(f_e^0)\cdot f_e^0 \\
&= L_0 - \sum_{e:\ell_e < (1+\delta')L_0} \ell_e(f_e^0)\cdot f_e^0
\end{aligned}$$

Since f^0 is at a (δ', ϵ')-equilibrium,

$$\begin{aligned}
\Phi(f^0) - \Phi(f') &\leq L_0 - (1-\delta')(1-\epsilon')L_0 \\
&\leq (\epsilon' + \delta')\, L_0 \\
&\leq 2\,d\,(\epsilon' + \delta')\,\Phi(f^0)\ . \qquad (4.11)
\end{aligned}$$

Since at most a total flow of v was moved within the phase and every edge had latency at least $(1-\delta)\,L_0$ at least once within the phase, there exists a flow vector $(\gamma_e)_{e\in E}$ of total volume at most v which can be added to f' such that the latency on every edge is at least $(1-\delta)\,L_0$. This flow vector only uses edges with $\ell_e(f_0) \leq (1-\delta)\,L_0$. For this flow $f'' = f' + \gamma$, clearly

$$\Phi(f') \leq \Phi(f'')\ . \qquad (4.12)$$

The total flow volume of f'' is $1+v$. Lemma 4.11 asserts, that the optimal potential $\tilde{\Phi}^*$ for this flow demand is at most $(1+3\,d\,v)\Phi^*$. Since by construction for all edges $e \in E$, $\ell_e(f''_e) \geq (1-\delta)L_0$ and for all edges with $f''_e > 0$, $\ell_e(f''_e) \leq (1+\delta')L_0$, Lemma 4.10 asserts that

$$\Phi(f'') \leq (1 + \mathcal{O}\left(d\,\delta\right))\ \tilde{\Phi}^* \leq (1 + \mathcal{O}\left(d\,\delta\right))\ (1 + 3\,d\,v)\,\Phi^*\ . \qquad (4.13)$$

Using Equations (4.11), (4.12), and (4.13),

$$\begin{aligned}
\Phi(f^0)\cdot(1 - 2\,d\,(\epsilon' + \delta')) &\leq \Phi(f') \\
&\leq \Phi(f'') \\
&\leq (1 + \mathcal{O}\left(d\,\delta + d\,v\right))\Phi^*\ .
\end{aligned}$$

Since $v = \mathcal{O}\left(\epsilon/d^2\right)$, this implies that, if c in the definition of $\delta = c\,\epsilon/d$ is chosen sufficiently small, f_0 is a $(1+\epsilon)$-approximation of the optimal potential and event (3) occurs. □

Theorem 4.13. *Consider a symmetric singleton instance Γ and an initial flow vector f_0. If $0 < \beta \leq \epsilon^2/(d^3 \cdot \ln(|\mathcal{P}|\,d/\epsilon))$, the ($\alpha$,$\beta$)-exploration-replication policy reaches a configuration f with potential $\Phi(f) \leq (1+\epsilon)\Phi^* + \epsilon\,\alpha$ in at most*

$$\mathcal{O}\left(\frac{d^{12}}{\epsilon^7}\cdot \ln^4\left(\frac{|E|}{\beta}\right)\cdot \ln\left(\frac{\Phi(f_0)+\alpha}{\Phi^*+\alpha}\right)\right)$$

rounds.

Proof. Again, by Fact 4.3, it is sufficient to prove the theorem for the case $\alpha = 0$ (see the first paragraph of the proof of Theorem 4.6). Consider a phase of length T as defined in Lemma 4.12. By Lemma 4.12, the phase terminates with one of the following events after at most T rounds:

(1) The population is no longer at a (δ',ϵ')-equilibrium. By Theorem 4.6, this can happen at most

$$T_1 = \mathcal{O}\left(\frac{d}{\epsilon' \delta'^2} \cdot \ln\left(\frac{\Phi(f^0)}{\Phi^*}\right)\right)$$

times.

(2) The potential decreases by at least $\Omega\left(\delta'/(d+1)\right) \cdot L_0$. This decreases the potential by a factor of at least $(1 - \Omega\left(\delta'/(d+1)\right))$. Therefore, this can happen at most

$$T_2 = \mathcal{O}\left(\frac{d}{\delta'} \cdot \ln\left(\frac{\Phi(f^0)}{\Phi^*}\right)\right)$$

times.

(3) The policy has reached a $(1+\epsilon)$-approximation of the potential.

Hence, after at most $T \cdot (T_1 + T_2)$ rounds, event (3) must occur. □

4.4.2.2 General Symmetric Games

We now consider the general symmetric case, i. e., $D \geq 2$. A lemma similar to Lemma 4.12 holds with modified definitions of δ, ϵ', δ', and T.

Lemma 4.14. *Consider a single-commodity instance Γ and the $(0, \beta)$-exploration-replication policy for $0 < \beta \leq \epsilon^4/(D^6\, d^5\, \ln(|\mathcal{P}|\, D\, d/\epsilon))$. For every $\epsilon > 0$ define the following constants:*

$$\delta = \frac{c\,\epsilon^2}{D^3\, d^2}\,, \quad \delta' = \epsilon' = c' \frac{\delta^2}{d \ln(\delta\, |\mathcal{P}|/\beta)}\,, \text{ and } \quad T = \frac{5\,d}{\lambda\,\delta} \ln\left(\frac{3\,\epsilon'\,d\,|\mathcal{P}|}{\lambda\,\delta\,\beta}\right),$$

where c and c' are positive constants independent of ϵ, β, and d to be defined in the proof. Then, in every phase consisting of T rounds starting with a flow vector f^0 at least one of the following events occurs:

(1) At least once in the phase the flow is not at a (δ',ϵ')-equilibrium.

(2) At the end of the phase, $\Phi(f) \leq \Phi(f^0) - \Omega\left(\frac{\delta'}{d}\right) L(f^0)$.

(3) $\Phi(f^0) \leq (1+\epsilon)\Phi^*$.

Proof. Consider a phase of length T. Throughout the proof, we write $L_0 = L(f^0)$. As in the proof of Lemma 4.12, assuming that events (1) and (2) do not occur, Lemma 4.9 guarantees that the average does not become smaller than $(1 - \delta/4)\, L_0$ if c' is chosen small enough. Again, the total amount of flow moved within the phase is at most $v = T \cdot \lambda\,(2\,\epsilon' + 2\,\delta' + \beta)/d$, and each path has latency $(1-\delta)L_0$ at least once within the phase. Due to the modified definition of δ, this time $v = \mathcal{O}\left(\epsilon^2/(D^3 d^3)\right)$.

Hence, there exists a flow vector with total volume v that can be added to the initial flow f^0 such that every path P with initial latency $\ell_P(f^0) \leq (1-\delta)\,L_0$ has latency at least $(1-\delta)\,L_0$. Unfortunately, we do not know which paths are used by this flow, but since every path may contribute to at most D edges, there exists a flow vector $(\gamma_P)_{P\in\mathcal{P}}$ with total volume $V = D\,v$ that only uses paths P with $\ell_P(f^0) \leq (1-\delta)\,L_0$ and increases the latency of no edge to above $(1-\delta)\,L_0$ to achieve the same goal. We fix such a flow vector γ and prove the following claim to establish the lemma.

Claim 4.15. *For the modified flow vector $\tilde{f} = f^0 + \gamma$ it holds that*

$$\Phi(\tilde{f}) \leq (1+\epsilon/2)\,\tilde{\Phi}^*$$

where $\tilde{\Phi}^*$ is defined as in Lemma 4.11.

The claim implies that

$$\begin{array}{rcll} \Phi(f^0) & \leq & \Phi(\tilde{f}) & \text{(monotonicity)} \\ & \leq & (1+\epsilon/2)\,\tilde{\Phi}^* & \text{(Claim 4.15)} \\ & \leq & (1+\epsilon/2)\,(1+3\,d\,V)\,\Phi^* & \text{(Lemma 4.11)} \\ & \leq & (1+\epsilon)\,\Phi^* \;\; , & \text{(definition of } V\text{)} \end{array}$$

which means that event (3) occurs.

Proof of Claim 4.15. The proof proceeds in two steps. First, we show that a significant increase of latency due to the additional flow affects only few agents. Then, we construct an instance in which the maximum latency of used paths is close to the minimum latency without changing the potential significantly and apply Lemma 4.10.

Consider the additional flow γ. Adding this flow may increase the latency of some paths significantly. Since, however, the flow on a resource will not be increased as soon as its latency has reached the value $(1-\delta)L_0$, we can be sure that we will increase the latency of no path to more than $D\,L_0$. In the following we will give a bound on the volume of agents utilising paths whose latencies are increased significantly. For any $e \in E$, let $\gamma_e = \sum_{P\ni e}\gamma_P$ and define $C = \sqrt{D^3\,\delta}/d = \mathcal{O}\left(\epsilon/d^2\right)$. For every resource e there are two cases.

Case 1. $\gamma_e \leq C\,f_e$. Then, $\tilde{f}_e = f_e + \gamma_e \leq f_e\,(1+C)$ implying that the latency on this resource does not increase by a factor of more than $1 + 2\,d\,C$.

Case 2. $\gamma_e > C\,f_e$. Since $\tilde{f}_e = f_e + \gamma_e \leq (1 + 1/C)\gamma_e$ and each path used by γ contributes to at most D edges, the total volume of agents utilising edges of this kind is at most

$$D \cdot \left(1 + \frac{1}{C}\right) \sum_{P \in \mathcal{P}} \gamma_P \leq \frac{2\,V\,D}{C} = \mathcal{O}\left(\frac{\epsilon}{D\,d}\right) \ .$$

In order for the latency of a path to grow from at most $\ell_P(f^0)$ to above $(1 + 2\,d\,C)\,\ell_P(f^0)$, it must contain at least one resource whose latency increases by a factor of at least $(1 + 2\,d\,C)$, i. e., a resource in case 2. Therefore, the total volume of agents utilising paths whose latencies are increased to above $(1 + 2\,d\,C)\,\ell_P(f^0)$ (but by at most $L_0\,D$) is bounded from above by $2\,V\,D/C$.

We modify the instance Γ to obtain an instance Γ' that reduces the latencies again down to their original values. First consider the paths

$$\mathcal{P}^1 = \{P \in \mathcal{P} \mid \ell_P(\tilde{f}) > (1 + 2\,d\,C)\,\ell_P(f^0)\} \ .$$

To every path $P \in \mathcal{P}^1$ we append a new resource with constant negative latency $\ell_P(f^0) - \ell_P(\tilde{f}) \geq -D\,L_0$. Denote the resulting path latency functions by ℓ'_P and note that $\ell'_P(\tilde{f}) = \ell_P(f^0)$ for all $P \in \mathcal{P}^1$ and $\ell'_P(\tilde{f}) = \ell_P(\tilde{f})$ for all $P \notin \mathcal{P}^1$. Due to our above considerations,

$$\Phi'(\tilde{f}) \geq \Phi(\tilde{f}) - L_0\,D \cdot \frac{2\,V\,D}{C} \ . \tag{4.14}$$

Now consider the paths

$$\mathcal{P}^2 = \{P \in \mathcal{P} \mid \ell'_P(\tilde{f}) \geq (1 + 3\,d\,C)\,L_0\} \ .$$

First, for any path $P \in \mathcal{P}^2 \cap \mathcal{P}^1$, $\ell_P(f^0) = \ell'_P(\tilde{f}) \geq (1 + 3\,d\,C)\,L_0$. Second, for any path $P \in \mathcal{P}^2 \setminus \mathcal{P}^1$ we know that

$$(1 + 3\,d\,C)\,L_0 \leq \ell'_P(\tilde{f}) = \ell_P(\tilde{f}) \leq (1 + 2\,d\,C)\,\ell_P(f^0)$$

implying $\ell_P(f^0) \geq (1 + d\,C)\,L_0$. In both cases $\ell_P(f^0) > (1 + \delta')\,L_0$. This implies that first, $\gamma_P = 0$, so $\tilde{f}_P = f^0_P$ and second, $\sum_{P \in \mathcal{P}^2} \tilde{f}_P \leq \epsilon'$ since f^0 is at a (δ',ϵ')-equilibrium. We transform Γ' once more into an instance Γ'' by appending a resource to any path $P \in \mathcal{P}^2$ with constant negative latency $L_0 - \ell'_P(\tilde{f})$. Then,

$$\begin{aligned}
\Phi'(\tilde{f}) - \Phi''(\tilde{f}) &\leq \sum_{P \in \mathcal{P}^2} \ell'_P(\tilde{f}) \cdot \tilde{f}_P \\
&\leq \sum_{P \in \mathcal{P}^2} (\ell_P(f_0) + D\,L_0) \cdot f^0_P \\
&\leq \epsilon'\,D\,L_0 + \sum_{P \in \mathcal{P}^2} \ell_P(f_0) \cdot f^0_P \\
&= \epsilon'\,D\,L_0 + L_0 - \sum_{P \notin \mathcal{P}^2} \ell_P(f_0) \cdot f^0_P \ .
\end{aligned}$$

Since f^0 is at a (δ',ϵ')-equilibrium,

$$\begin{aligned}
\Phi'(\tilde{f}) - \Phi''(\tilde{f}) &\leq \epsilon' D\, L_0 + L_0 - (1-\epsilon')(1-\delta')\, L_0 \\
&\leq \epsilon' D\, L_0 + (\epsilon' + \delta') L_0 \\
&\leq 2\,\epsilon' D\, L_0 && (4.15)
\end{aligned}$$

since $D \geq 2$. We now have an instance Γ'' and a modified flow $\tilde{f}$ with the property that the minimum latency is $(1-\delta)L_0$ and the maximum latency of used paths is $(1 + 3\,d\,C)L_0$. Since $\delta \leq 3\,d\,C$, Lemma 4.10 states, that $\tilde{f}$ is a $\left(1 + \mathcal{O}\left(d^2\,C\right)\right)$-approximation of the optimal potential of Γ''. Consequently,

$$\begin{aligned}
\Phi''^* + \mathcal{O}\left(d^2\,C\right)\,\Phi''^* &\geq \Phi''(\tilde{f}) \\
&\geq \Phi'(\tilde{f}) - 2\,\epsilon' D\, L_0 && \text{(Eq. (4.15))} \\
&\geq \Phi(\tilde{f}) - 2\,\epsilon' D\, L_0 - \tfrac{2\,V\,D^2}{C}\, L_0 && \text{(Eq. (4.14))} \\
&\geq \Phi(\tilde{f}) - \tfrac{3\,V\,D^2}{C} L(\tilde{f}) && (L(\tilde{f}) \geq L_0) \\
&\geq \Phi(\tilde{f})\left(1 - \tfrac{6\,d\,V\,D^2}{C}\right) \;. && \text{(Fact 4.2)}
\end{aligned}$$

Hence, we have

$$\begin{aligned}
\Phi(\tilde{f}) &\leq \Phi''^*\left(1 + \mathcal{O}\left(d^2\,C + \tfrac{d\,V\,D^2}{C}\right)\right) \\
&\leq \Phi''^*\,(1 + \epsilon/2) && \text{(Def. of } V \text{ and } C)
\end{aligned}$$

if the constant c in the definition of δ is chosen sufficiently small. □

The proof of the claim completes the proof of the lemma. □

Theorem 4.16. *Consider a symmetric instance Γ and an initial flow vector f_0. If $0 < \beta \leq \epsilon^4/(D^6\,d^5\,\ln(|\mathcal{P}|\,D\,d/\epsilon))$, the ($\alpha$,$\beta$)-exploration-replication policy generates a configuration with potential $\Phi \leq (1+\epsilon)\Phi^* + \epsilon\,\alpha$ in at most*

$$\operatorname{poly}\left(d, \frac{1}{\epsilon}, D\right) \ln^4\left(\frac{|E|}{\beta}\right) \ln\left(\frac{\Phi(f_0) + \alpha}{\Phi^* + \alpha}\right)$$

rounds.

The proof is identical with the proof of Theorem 4.13 with the modified definitions of δ, ϵ', δ', and T. In addition, we use that $|\mathcal{P}| = \mathcal{O}\left(|E|^D\right)$. Substituting our bound on β as specified by Equation (4.1), this yields the bound as presented in Section 4.1.1.

4.5 Convergence Time for Multi-Commodity Instances

For the asymmetric, or multi-commodity, case, we generalise the bicriteria definition of approximate equilibria in the following manner.

Definition 4.6 ((δ,ϵ)-equilibrium). *For a flow vector $f \in \mathcal{F}_\Gamma$ and commodity $i \in [k]$, let $\mathcal{P}_i^+(\delta) = \{P \in \mathcal{P}_i \mid \ell_P(f) \geq L_i(f) + \delta L(f)\}$ denote the set of* δ-expensive paths *and let $\mathcal{P}_i^-(\delta) = \{P \in \mathcal{P}_i \mid \ell_P(f) \leq L_i(f) - \delta L(f)\}$ denote the set of* δ-cheap paths. *The population f is at a (δ,ϵ)-equilibrium iff*

$$\sum_{i\in[k]} \sum_{P\in\mathcal{P}_i^+(\delta)\cup\mathcal{P}_i^-(\delta)} f_P \leq \epsilon \ .$$

The analysis of the time of convergence is slightly more involved than in the symmetric case.

Theorem 4.17. *Consider an asymmetric instance Γ and an initial flow vector $f_0 \in \mathcal{F}_\Gamma$. For the ($\alpha$,$\beta$)-exploration-replication policy, the number of rounds in which the population vector is not at a (δ,ϵ)-equilibrium w. r. t. $\Gamma^{+\alpha}$ (Definition 4.2) is bounded from above by*

$$\mathcal{O}\left(\frac{d}{\epsilon^2\,\delta^2}\log\left(\frac{\Phi(f_0)+\alpha}{\Phi^*+\alpha}\right)\right) \ .$$

In particular, this bound holds for $\alpha = \beta = 0$ (and hence $\Gamma^{+\alpha} = \Gamma$).

Proof. As in the proof of Theorem 4.6, it is sufficient to consider the case $\alpha = 0$. Again, we estimate the virtual potential gain $\mathcal{V}$. Consider a commodity $i \in [k]$. There are two cases.

Case 1. $L_i > L \cdot 2/\epsilon$. By Markov's inequality, at most a volume of $\epsilon/2$ agents may belong to commodities of this type. We ignore the potential gain of these agents.

Case 2. $L_i \leq L \cdot 2/\epsilon$. The remaining $1 - \epsilon/2$ agents belong to commodities of this type. As long as we are not at a (δ,ϵ)-equilibrium, there must be at least a volume of ϵ agents utilising δ-expensive or δ-cheap paths. Since at most $\epsilon/2$ of them are in case 1, at least $\epsilon/2$ of them are in this case.

Throughout the proof we only consider agents of the second case. For these agents, the proof is an extension of the proof of Theorem 4.6. Intuitively, for the proof of Theorem 4.17 we have to adapt only a single argument: In the single-commodity case, an expensive path is by δL more expensive than the average, and agents on expensive paths that sample an average latency path migrate to this path with a probability of almost δ. In the multi-commodity case, however, the average latency of a commodity may be larger than the overall average latency, but for the commodities we consider, $L_i \leq L \cdot 2/\epsilon$. Hence, a deviation from L_i by δL still implies a migration probability of almost $\delta \epsilon/2$. This is by a factor of $\epsilon/2$ smaller than the migration probabilities in the proof of Theorem 4.6.

Fix a commodity $i \in [k]$ and let x_i denote the volume of agents utilising δ-cheap or δ-expensive paths, whichever is larger.

Case 1. The volume of agents in commodity i utilising δ-expensive paths is larger than the volume of agents utilising δ-cheap paths, i. e., x_i is the volume of agents utilising δ-expensive paths.

Case 1a. In commodity i, at least a volume of $x_i \geq r_i/2$ agents utilises δ-expensive paths. Let $y > 0$ denote the volume of agents utilising paths with latency at most L_i. Consider a random agent that is sampled by proportional sampling and let the random variable Y denote its current latency. Let $\Delta = \mathbb{E}\left[L_i - Y \mid Y \leq L_i\right]$ denote the conditional expectation of the difference between Y and L_i. We have

$$\Delta = L_i - \frac{r_i}{y}\left(\sum_{P \in \mathcal{P}_i : \ell_P \leq L_i} \frac{f_P}{r_i} \ell_P\right) .$$

Furthermore, by definition of x_i,

$$\sum_{P \in \mathcal{P}_i : \ell_P > L_i + \delta L} \frac{f_P}{r_i} \ell_P \geq \frac{\delta\, x_i}{r_i}\, L + \frac{x_i}{r_i}\, L_i .$$

Finally, clearly

$$\sum_{P \in \mathcal{P}_i : L < \ell_P \leq L_i + \delta L} \frac{f_P}{r_i} \ell_P \geq \frac{r_i - x_i - y}{r_i} \cdot L_i .$$

Substituting the last three equations into the definition of L_i yields

$$\begin{aligned} L_i &= \sum_{P \in \mathcal{P}_i : \ell_P \leq L_i} \frac{f_P}{r_i} \ell_P + \sum_{P \in \mathcal{P}_i : L_i < \ell_P \leq L_i + \delta L} \frac{f_P}{r_i} \ell_P + \sum_{P \in \mathcal{P}_i : \ell_P > L_i + \delta L} \frac{f_P}{r_i} \ell_P \\ &\geq \frac{y \cdot (L_i - \Delta) + (r_i - x_i - y) \cdot L_i + \delta\, x_i\, L + x_i\, L_i}{r_i} \\ &= \frac{r_i\, L_i - y\, \Delta + \delta\, x_i\, L}{r_i} \end{aligned}$$

implying $\Delta \geq (\delta\, x_i / y)\, L$. All agents with latency at least $L_i + \delta L$ that sample a path with latency at most L_i migrate to the new path with probability at least $\delta\,\epsilon/(2\,d\,(1+\delta))$ as discussed above. The probability to sample such a path is at least $(1-\beta)\,y/r_i$. Their expected latency gain which is equivalent to their infinitesimal contribution to the virtual potential gain is at least Δ. In total there are $x_i \geq r_i/2$ such agents. Altogether, the virtual potential gain of commodity i is

$$\mathcal{V}_i \leq -x_i \cdot \frac{\lambda\,(1-\beta)\,y\,\delta\,\epsilon}{2\,d\,r_i\,(1+\delta)} \cdot \Delta \leq -\frac{\lambda\,(1-\beta)\,x_i^2\,\delta^2\,\epsilon}{2\,d\,r_i\,(1+\delta)} L \leq -r_i\,\frac{\lambda\,\delta^2\,\epsilon}{16\,d} L \ .$$

Case 1b. The volume of agents in commodity i utilising δ-expensive paths in commodity i is $x_i \leq r_i/2$. Then, for this commodity, the virtual potential gain of the agents leaving δ-expensive paths is at least

$$\mathcal{V}_i \quad \leq \quad -\frac{\lambda}{d} \sum_{P\in\mathcal{P}_i^+} f_P \sum_{Q\notin\mathcal{P}_i^+} \left((1-\beta)\frac{f_Q}{r_i} + \frac{\beta}{|\mathcal{P}_i|}\right) \frac{(\ell_P-\ell_Q)^2}{\ell_P} \ .$$

Omitting the term $\beta/|\mathcal{P}_i|$, substituting $\ell_P \geq L_i + \delta L$ and $(1-\beta) \geq 1/2$ and applying Jensen's inequality to the last sum yields

$$\begin{aligned}
\mathcal{V}_i \quad &\leq \quad -\frac{\lambda}{2\,d\,(L_i+\delta\,L)} \sum_{P\in\mathcal{P}_i^+} f_P \left(\sum_{Q\notin\mathcal{P}_i^+} \frac{f_Q}{r_i}(L_i+\delta L-\ell_Q)\right)^2 \frac{1}{\sum_{Q\notin\mathcal{P}_i^+} f_Q/r_i} \\
&\leq \quad -\frac{\lambda}{2\,d\,(L_i+\delta\,L)} \sum_{P\in\mathcal{P}_i^+} f_P \left((L_i+\delta L)\sum_{Q\notin\mathcal{P}_i^+} \frac{f_Q}{r_i} - \sum_{Q\notin\mathcal{P}_i^+} \frac{f_Q}{r_i}\ell_Q\right)^2 \ .
\end{aligned}$$

Note that $\sum_{Q\notin\mathcal{P}_i^+}(f_Q/r_i)\ell_Q \leq L_i \sum_{Q\notin\mathcal{P}_i^+}(f_Q/r_i)$ since the sum omits the terms of the expensive paths $Q \in \mathcal{P}_i^+$. Using our assumption $L_i \leq 2 \cdot L/\epsilon$,

$$\begin{aligned}
\mathcal{V}_i \quad &\leq \quad -\frac{\lambda}{2\,d\,(2/\epsilon+\delta)\,L} \sum_{P\in\mathcal{P}_i^+} f_P \left(\delta\,L \sum_{Q\notin\mathcal{P}_i^+} \frac{f_Q}{r_i}\right)^2 \\
&\leq \quad -\frac{\lambda\,\delta^2\,L^2}{2\,d\,(2/\epsilon+\delta)\,L} \sum_{P\in\mathcal{P}_i^+} f_P \left(\sum_{Q\notin\mathcal{P}_i^+} \frac{f_Q}{r_i}\right)^2 \\
&\leq \quad -x_i \frac{\lambda\,\epsilon\,\delta^2}{24\,d} L \ .
\end{aligned}$$

In the last inequality we used our assumption that $\sum_{Q\notin\mathcal{P}_i^+} f_Q \geq r_i/2$.

Case 2. The volume of agents in commodity i utilising δ-cheap paths is larger than the volume of agents utilising δ-expensive paths, i. e., x_i is the volume of agents utilising δ-cheap paths.

Case 2a. In commodity i, at least a volume of $x_i \geq r_i/2$ utilises δ-cheap paths. Let $y > 0$ denote the volume of agents utilising paths with latency at least L_i. Consider a random agent and let the random variable Y denote its current latency. Let $\Delta = \mathbb{E}\left[L_i - Y \mid Y \geq L_i\right]$ denote the conditional expectation of the difference between L_i and Y. Then,

$$\begin{aligned}
\Delta &= \frac{r_i}{y}\left(\sum_{P \in \mathcal{P}_i : \ell_P \geq L_i} \frac{f_P}{r_i}\ell_P\right) - L_i \ , \\
\sum_{P \in \mathcal{P}_i : \ell_P < L_i - \delta L} \frac{f_P}{r_i}\ell_P &\leq \frac{\delta\, x_i}{r_i}\, L - \frac{x_i}{r_i}\, L_i \ \text{, and} \\
\sum_{P \in \mathcal{P}_i : L - \delta L \leq \ell_P < L_i} \frac{f_P}{r_i}\ell_P &\leq \frac{r_i - x_i - y}{r_i} \cdot L_i \ .
\end{aligned}$$

Substituting this into the definition of L_i,

$$\begin{aligned}
L_i &\leq \frac{y \cdot (L_i + \Delta) + (r_i - x_i - y) \cdot L_i - \delta\, x_i\, L + x_i\, L_i}{r_i} \\
&= \frac{r_i\, L_i + y\, \Delta - \delta\, x_i\, L}{r_i}
\end{aligned}$$

implying $\Delta \geq (\delta\, x_i/y)\, L$. All agents with latency at least L_i that sample a path with latency at most $L_i - \delta\, L$ migrate to the new path with probability at least $\delta\, \epsilon/(2\, d)$. The probability to sample such a path is at least $(1-\beta)\, x_i/r_i$. Their expected latency gain which is equivalent to their infinitesimal contribution to the virtual potential gain is at least Δ. In total there are y such agents. Again, the virtual potential gain of commodity i is

$$\mathcal{V}_i \leq -y \cdot \frac{\lambda\,(1-\beta)\, x_i\, \delta\, \epsilon}{2\, d\, r_i} \cdot \Delta \leq -r_i\, \frac{\lambda\, \delta^2\, \epsilon}{16\, d} L \ .$$

Case 2b. In commodity i, at most a volume of $r_i/2$ agents utilise δ-cheap paths. We can treat the virtual potential gain of agents moving towards δ-cheap paths in a way similar to case 1b.

$$\begin{aligned}
\mathcal{V}_i &\leq -\frac{\lambda}{d} \sum_{P \in \mathcal{P}_i^-} \left((1-\beta)\frac{f_P}{r_i} + \frac{\beta}{|\mathcal{P}_i|}\right) \sum_{Q \notin \mathcal{P}_i^-} f_Q \frac{(\ell_Q - \ell_P)^2}{\ell_Q} \\
&\leq -\frac{\lambda}{d} \sum_{P \in \mathcal{P}_i^-} ((1-\beta) f_P) \sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i}\, \ell_Q \left(\frac{\ell_Q - \ell_P}{\ell_Q}\right)^2 \ .
\end{aligned}$$

Now, we use that $\ell_P \leq L_i - \delta L$ and $L_i \leq 2\, L/\epsilon$ and apply Jensen's inequality.

$$
\begin{aligned}
\mathcal{V} &\leq -\frac{\lambda}{3\,d} \sum_{P \in \mathcal{P}_i^-} f_P \left(\sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} \ell_Q \frac{\ell_Q - \ell_P}{\ell_Q} \right)^2 \frac{1}{\sum_{Q \notin \mathcal{P}_i^-} (f_Q / r_i) \ell_Q} \\
&\leq -\frac{\lambda}{3\,d\,L_i} \sum_{P \in \mathcal{P}_i^-} f_P \left(\sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} (\ell_Q - L_i + \delta L) \right)^2 \\
&\leq -\frac{\lambda\,\epsilon}{3\,d\,L} \sum_{P \in \mathcal{P}_i^-} f_P \left(\sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} \ell_Q + (\delta L - L_i) \sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} \right)^2 .
\end{aligned}
$$

Note that $\sum_{Q \notin \mathcal{P}_i^-} (f_Q / r_i)\, \ell_Q \geq L_i \sum_{Q \notin \mathcal{P}_i^-} (f_Q / r_i)$ since the sum omits the terms of the paths $Q \in \mathcal{P}_i^-$ with low latency. Hence,

$$
\begin{aligned}
\mathcal{V} &\leq -\frac{\lambda\,\epsilon}{3\,d\,L} \sum_{P \in \mathcal{P}_i^-} f_P \left(\delta\,L \sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} \right)^2 \\
&\leq -\frac{\lambda\,\epsilon\,\delta^2\,L^2}{3\,d\,L} \sum_{P \in \mathcal{P}_i^-} f_P \left(\sum_{Q \notin \mathcal{P}_i^-} \frac{f_Q}{r_i} \right)^2 \\
&\leq -x_i \cdot \frac{\lambda\,\epsilon\,(1 - 1/2)^2\,\delta^2}{3\,d} L \; .
\end{aligned}
$$

Finally, summing up over all commodities, the virtual potential gain is

$$
\begin{aligned}
\mathcal{V} &\leq -\sum_{i \in [k]} \mathcal{V}_i \\
&\leq -\sum_{i \in [k]} \min \left\{ r_i \frac{\lambda\,\delta^2\,\epsilon\,L}{16\,d}, x_i \frac{\lambda\,\epsilon\,\delta^2\,L}{24\,d} \right\} \\
&\leq -\frac{\lambda\,\epsilon\,\delta^2\,L}{24\,d} \sum_{i \in [k]} x_i \\
&\leq -\frac{\lambda\,\epsilon^2\,\delta^2\,L}{96\,d} \; .
\end{aligned}
$$

Here, the minimum is over the two subcases (a) and (b), and the last inequality uses that $\sum_{i \in [k]} x_i \geq \epsilon/4$. By Lemma 4.4 the potential decreases by at least half of this in one round. Given this rate at which the potential decreases, we obtain the desired bound as in the proof of Theorem 4.6. □

4.6 Lower Bounds

The results obtained in the preceding sections rely on bounded elasticity of the latency functions and using proportional sampling. In the following, we present two lower bounds showing that adaptive sampling is necessary to obtain polynomial time bounds and that the dependence on elasticity cannot be avoided.

4.6.1 Elasticity is Necessary

In this section we show that elasticity is the relevant parameter in our analysis. To that end, we consider arbitrary rerouting policies that avoid oscillation for a given class of latency functions $\mathcal{L}$. We formalise this by requiring that for any feasible population, the policy may not increase the potential.

Definition 4.7 ($\mathcal{L}$-monotone)**.** *Fix a network $G = (V, E)$ and a set of commodities $\mathcal{C}$. For any class of latency functions $\mathcal{L}$, a rerouting policy is* $\mathcal{L}$-monotone *if for all instances $\Gamma = (G, \mathcal{C}, (\ell_e)_{e \in E})$ with $(\ell_e) \in \mathcal{L}^E$ and all flow vectors $f \in \mathcal{F}_\Gamma$ it holds that for the flow $f' \in \mathcal{F}_\Gamma$ generated by the policy from f, $\Phi(f') \leq \Phi(f)$.*

Lemma 4.4 ensures that the (α,β)-exploration-replication policy is $\mathcal{L}$-monotone if $\mathcal{L}$ has elasticity d. We provide a lower bound showing that for a class $\mathcal{L}$ of latency functions the elasticity d of $\mathcal{L}$ is a lower bound for the time of convergence towards approximate equilibria if the rerouting policy is $\mathcal{L}$-monotone.

Theorem 4.18. *For every d, there exists a class $\mathcal{L}$ of latency functions with elasticity d together with an initial flow vector f_0, such that any $\mathcal{L}$-monotone rerouting policy requires $\Omega(d/\epsilon)$ rounds in order to obtain a $(1 + \epsilon)$ approximation of the optimum potential.*

Proof. We choose positive constants $a, b > 0$ and an arbitrary latency function $\ell : [0, 1] \mapsto \mathbb{R}_{\geq 0}$ that has elasticity d in the interval $[a, b]$. (As an example, simply choose $\ell(x) = x^d$, $a = 1/2$, and $b = 1$.) We now define a set $\mathcal{L}$ of latency functions which have elasticity at most d. For a parameter $u \in [0, b - a]$, we define the function $\ell_u(x) = \min\{\ell(x + u), \ell(b)\}$ obtained from ℓ by shifting it to the left after cutting it off at b. Note that $\ell_0 = \ell$ has elasticity d and for $u > 0$, the elasticity of ℓ_u is less than d. Furthermore, we define the constant latency functions $\ell_1(x) = \ell_1 = \ell(b)$ and $\ell_2(x) = \ell_2 = \ell(b)/(1 + \epsilon/a)$. The class of latency functions we consider is

$$\mathcal{L} := \{\ell_u \mid u \in [0, b - a]\} \cup \{\ell_1, \ell_2\} \ .$$

We define an instance consisting of two parallel links with the constant latency functions ℓ_1 and ℓ_2, respectively. We denote the flows on the respective resources by f_1 and $f_2 = 1 - f_1$. Note that the optimal potential is $\Phi^* = \ell_2$ and for an approximate equilibrium with $\Phi(f) \leq (1+\epsilon)\Phi^*$ we must have $f_1 \leq a$.

Now consider any $\mathcal{L}$-monotone rerouting policy and let $f(t)_{t\in\mathbb{N}}$ denote the sequence of flow vectors generated by this policy. As a starting configuration we choose $f_1(0) = b$ and $f_2(0) = 1-b$. Given this sequence, let $\mathcal{L}(0) = \mathcal{L}$ and for $t > 0$ let

$$\mathcal{L}(t) = \{\ell \in \mathcal{L}(t-1) \mid \ell(f_1(t-1)) = \ell_1\} \ .$$

These are the latency functions that are compatible with the flows and latencies observed on resource 1 by the policy in the first t rounds. Note that the monotonicity of the policy ensures that $f_1(t)$ does not increases with t. Furthermore, for any $u \in [0, b-a]$ and $x \geq b-u$ we have $\ell_u(x) = \ell_1$. In particular, $\ell_u(f_1(t)) = \ell_1$ for all $u \in [b - f_1(t), b-a]$ implying that ℓ_u is (and was in previous rounds) compatible with the observed latencies on resource 1 and hence $\ell_u \in \mathcal{L}(t)$.

At time t, the rerouting policy must guarantee that it does not increase the potential for all possible latency functions in $\mathcal{L}(t)$. In particular, this must hold for the case that the latency function of the first resource is $\ell_{b-f_1(t)}$. More precisely, the policy must ensure that

$$\ell_{b-f_1(t)}(f_1(t+1)) \geq \ell_2 \cdot (1 - 3\cdot\epsilon/a) \ .$$

How much flow can be shifted from resource 1 to resource 2 without violating this inequality? Note that the left derivative of ℓ_u at $b-u$ equals the derivative of ℓ at b and since ℓ has elasticity d on $[a,b]$, it is at least $d \cdot \ell_1/b$. In particular, $\ell'_{b-f_1(t)}(f_1(t) - \delta) \geq d \cdot \ell_1/(2\,b)$ for small δ as an implication of Fact 4.1. Hence, in order to be sure that the potential does not increase, the policy must ensure that for the amount ρ of flow shifted from the first to the second path it holds that

$$\rho \cdot \frac{d\cdot\ell_1}{2\cdot b} \leq \frac{3\cdot\ell_2\cdot\epsilon}{a} \ ,$$

or, equivalently, $\rho \leq 6\,\epsilon/d \cdot (b/a)$. Since b and a are constants, this implies, that

$$f_1(t+1) \geq f_1(t) - \mathcal{O}(\epsilon/d) \geq f_1(0) - \mathcal{O}(t\,\epsilon/d) \ .$$

This implies that $f_1(t) \geq a$ as long as $t \leq c \cdot d/\epsilon$ for a suitable constant $c > 0$, our desired bound. □

Let us remark, that the instance presented in the proof can be used to show that the same bound holds for the time to reach a bicriterial (δ, c)-equilibrium with $\delta = \Omega(\epsilon)$ and constant $c > 0$.

Furthermore, note that this theorem does not require the policy to be Markovian. Furthermore, the proof shows that the theorem actually holds for any set of latency functions $\mathcal{L}$ containing at least one function ℓ that has elasticity d on an interval of constant width plus all functions obtained from this function by shifting it to the left and cutting it off at a particular value. Such a left-shift has an intuitive motivation. We can interpret it as a base load on the given edge.

4.6.2 Sampling with Static Probabilities is Slow

The following theorem shows that every rerouting policy that samples with static probabilities that are independent of the latency functions needs at least $\Omega(|P|)$ rounds to approach a Wardrop equilibrium. We will formalise the notion of static sampling probabilities in the following way. If path P is sampled with static probability σ_P, at most $(1 - f_P)\,\sigma_P$ agents may migrate towards P in one round since $(1 - f_P)$ agents utilise other paths. We say that a rerouting policy has *static sampling probabilities* (for a set of paths $\mathcal{P}$) denoted by $(\sigma_P)_{P\in\mathcal{P}}$ with $\sum_{P\in\mathcal{P}_i} \sigma_P = 1$ for all $i \in [k]$, if for every feasible flow vector $f \in \mathcal{F}_\Gamma$, every commodity $i \in [k]$, and every path $P \in \mathcal{P}_i$ it holds that the total volume of flow that the policy shifts to path P in one round is bounded from above by $\sigma_P(r_i - f_P)$.

Theorem 4.19. *For every m, there exist a graph with $|E| = m$ edges and $|\mathcal{P}| = 2^{m/4}$ paths such that for every rerouting policy with static sampling probabilities for $\mathcal{P}$ there exist a symmetric instance on G with constant latency functions and an initial flow vector f_0 such that the rerouting policy needs at least $\Omega(|\mathcal{P}|\log(1/\epsilon))$ rounds to reach a $(1+\epsilon)$-approximation of the optimal potential.*

Proof. Consider a network $G = (V, E)$ with vertices

$$V = \{s, t, v_1, w_1, \ldots, v_n, w_n\}$$

and an edge set consisting of the edges (s, v_1), (s, w_1), (v_n, t), (w_n, t), and, for $1 \le i < n$, (v_i, v_{i+1}), (v_i, w_{i+1}), (w_i, v_{i+1}), and (w_i, w_{i+1}). This network has $m = 4\,n$ resources and $|\mathcal{P}| = 2^n = 2^{m/4}$ paths.

Let $(\sigma_P)_{P\in\mathcal{P}}$ be the vector of sampling probabilities chosen by the rerouting policy for G and let $\tilde{P}$ be the path that minimises σ_P. Then, $\sigma_{\tilde{P}} \le 1/|\mathcal{P}|$.

We now define constant latency functions on G. First, we define $\ell_e = 0$ for any $e \in \{(s, v_1), (s, w_1), (v_n, t), (w_n, t)\}$. For every resource $e \in \tilde{P}$ with $s, t \notin e$, let $\ell_e = 1$. For all other resources $e' \notin \tilde{P}$, let $\ell_{e'} = n$. Hence, path $\tilde{P}$ is the unique optimal path with constant latency $n-1$, and $\Phi^* = n - 1$. All other paths have latency at least $2\,n - 2$. In order to reach a $(1+\epsilon)$ approximation of Φ^*, we must have

$$f_{\tilde{P}} \cdot (n-1) + (1 - f_{\tilde{P}}) \cdot (2\,n - 2) \le (1+\epsilon)\Phi^* \ ,$$

or, equivalently, $f_{\tilde{P}} \geq 1 - \epsilon$. By our choice of $\tilde{P}$, we have

$$f_{\tilde{P}}(t+1) \leq f_{\tilde{P}}(t) + (1 - f_{\tilde{P}}(t))\sigma_{\tilde{P}} \leq f_{\tilde{P}}(t) + \frac{1 - f_{\tilde{P}}(t)}{|\mathcal{P}|}$$

or, writing $\bar{f}_{\tilde{P}} = 1 - f_{\tilde{P}}$,

$$\bar{f}_{\tilde{P}}(t+1) \geq \bar{f}_{\tilde{P}}(t) - \frac{\bar{f}_{\tilde{P}}(t)}{|\mathcal{P}|} = \bar{f}_{\tilde{P}}(t)\left(1 - \frac{1}{|\mathcal{P}|}\right) .$$

Therefore, $\bar{f}_{\tilde{P}}(t) \geq \bar{f}_{\tilde{P}}(0) \cdot (1 - 1/|\mathcal{P}|)^t$ implying that it takes at least $\Omega(|\mathcal{P}| \cdot \log(1/\epsilon))$ rounds until $\bar{f}_{\tilde{P}}(t) < \epsilon$ and hence $f_{\tilde{P}} \geq 1 - \epsilon$ if we choose $f_{\tilde{P}}(0) > 0$ constant. □

4.7 Summary

We have introduced the (α,β)-exploration-replication policy which uses a hybrid sampling rule. It carefully explores the strategy space by using uniform sampling and boosts good strategies by proportional sampling. We have shown that this policy converges quickly towards Wardrop equilibria in a round-based model which takes into account possible oscillation effects due to the staleness of information. We have identified the *elasticity* of the latency functions as the crucial parameter limiting the speed of convergence. All of our bounds on the time of convergence towards approximate equilibria depend polynomially on this parameter. In particular, our bounds are polynomial in the encoding length of an instance if, e. g., latency functions are given as positive polynomials of arbitrary degree in coefficient representation. We consider this a significant improvement over earlier results regarding related models and objectives that rely on the maximum slope of the latency functions as the limiting parameter, thus having a pseudopolynomial flavour.

We have considered two different definitions of approximate equilibria. First, we have shown that the (α,β)-exploration-replication policy reaches a bicriterial approximate equilibrium in which almost all agents deviate by no more than a constant fraction of the overall average latency from the average latency of their commodity within a time that is polynomial in the approximation parameters but is almost independent of the network size. Since approximate equilibria in this bicriterial sense are transient, we have considered a second definition of approximate equilibria. For the symmetric case we have shown that our dynamics reaches a population which deviates by no more than a factor of $(1 + \epsilon)$ from the optimal potential in polynomial time. In this sense, our exploration-replication policy behaves like an FPAS.

Our analysis has shown that the elasticity and the proportional sampling rule are tightly coupled. Overshooting effects can be bounded in terms of the elasticity merely by the fact that we use proportional sampling which causes proportional growth rates of the edge flows. Two lower bounds show that an adaptive sampling rule is actually necessary to obtain polynomial time bounds and that any oscillation avoiding policy needs convergence time $\Omega(d)$.

The main question left open by this chapter is certainly whether a similar polynomial time bound can be obtained for the time to reach $(1+\epsilon)$-approximation of the optimal potential in the asymmetric case. It is not entirely obvious whether or not this is possible, since computing equilibria for the corresponding discrete congestion games is known to be PLS-complete [1, 29] in the asymmetric case.

CHAPTER 5

REPLEX – Dynamic Traffic Engineering

The positive results of the preceding chapters on load-adaptive rerouting policies encourage the design of dynamic traffic engineering protocols that are applicable in practise. Whereas our theoretical model encompasses one of the key problems arising in practise, namely the issue of stale information, real-world applications face several other problems. Traffic demands may be subject to unpredictable changes, traffic is bursty, and latency cannot be simply "observed" as assumed so far, but must be measured reliably. In view of these obstacles, today's traffic engineering protocols operate mainly off-line on the time scale of several hours or days and thus cannot react timely to effects like flash crowds or sudden changes of the network topology, e. g., due to link failures or BGP reroutes. We propose a dynamic traffic engineering protocol, REPLEX, based on the (α, β)-exploration-replication policy introduced in Chapter 4. Our protocol splits traffic that is bound for the same destination at intermediate routers and readjusts the splitting ratios at intervals of a few seconds. We perform simulations employing realistic workload generators on artificial topologies to inspect characteristic convergence properties and on a real topology to study its benefits on the overall performance. It turns out that our protocol converges within few minutes and improves the overall throughput significantly.

5.1 Introduction

We consider an autonomous system (AS) from the point of view of an Internet service provider (ISP) striving to maximise the benefit of its network. Whereas the topology of the network is fixed, the routing within the AS should be optimised for the traffic requirements between pairs of routers located at the border of the AS, so-called ingress and egress nodes. Typical optimisation goals are maximising the throughput of the network or minimising the maximum link load. In general, this task is referred to as *traffic engineering (TE)*.

We present a dynamic traffic engineering protocol, REPLEX, operated by routers which may or may not coordinate their actions. It assigns weights to routes (i. e., to outgoing edges) and continuously readjusts them in a fashion similar to the replication-exploration policy. Packets are then distributed evenly according to these weights. To that end, we need to assume that, for any given destination, a router may select between several next-hop routers. These alternative paths can be provided by any underlying routing protocol, e. g., by OSPF or MPLS multipaths. One advantage of our protocol is the fact that it is independent of the protocol which provides these alternatives. It can operate even across AS boundaries using BGP/MIRO multipaths [85]. Routers that apply our protocol may exchange traffic information at regular intervals. Since we aggregate information about several routes and transfer this information only to adjacent routers, the communication overhead generated by our protocol is very moderate.

How the weights are actually adjusted throughout the process depends on the optimisation goal pursued. Different applications pose different requirements on the performance of the connections. Typical requirements are low latency, high bit rate, low error rate, low packet loss, or a combination of these. Whereas our theoretical results focused on latency as a measure of performance, the protocol presented in this chapter is generic with respect to the measure of performance. We merely assume that the performance of a path can be recursively expressed as a function of the performance of the first link and the performance of the remaining path section. To mention two examples, latency can be summed up along a path, and the bit rate of a path can be computed by taking the minimum bit rate over all edges.

We evaluate our protocol using artificial topologies as well as a real topology from the Rocketfuel [3, 76] dataset. Since traffic engineering protocols may induce feedback to underlying protocols like TCP and may thus influence the traffic demands themselves, a realistic simulation must include a realistic workload generator. We incorporate such a generator [81] based on realistic statistical properties of Web traffic into our simulation framework. Our simulations are performed at packet level and include a full TCP implementation. Simulations show that our protocol converges within less than three minutes only and significantly increases the performance of the network within this time. A less significant improvement can still be observed if the communication between the routers is disabled.

Let us remark that our protocol aims at balancing the performance of different routes rather than optimising the overall performance. For certain performance measures like link utilisation, this is also a natural goal pursued by an ISP.

5.1.1 Related Work

The first approaches of on-line traffic engineering protocols date back to the times of the early ARPANET [47]. These protocols, however, tend to exhibit oscillating behaviour due to several reasons. First, routing decisions of individual routers may interact with each other and, second, decisions of the traffic engineering protocol may feed back to the traffic demands themselves: The TCP protocol which accounts for the largest fraction of today's Internet trafficlimits congestion by adjusting the injection rate at end hosts whenever a packet is dropped or packets are reordered, possibly as a result of a rerouting decision made by the TE protocol.

For these reasons, traffic engineering protocols applied in practise today are mainly off-line protocols. These collect traffic measurements over a certain period of time, centrally compute new routing parameters, and update routing tables based on these parameters. Typically, this process is re-run every few hours or days. One natural approach for controlling the routing paths is by readjusting OSPF weights [36, 37, 79]. The OSPF routing protocol determines routing paths by performing a weighted shortest path first computation. In [39] it is even proposed to adapt these weights dynamically. Any change of OSPF weights, however, requires the intradomain routing protocol to converge. Whereas traffic engineering extensions to OSPF have also been proposed [45], these merely describe how routers can add performance information relevant to a TE protocol to the transmitted messages. It does not address the question of how this information can actually be used for traffic engineering.

To a certain extent, off-line protocols can take into account events like link failures and compute alternatives off-line [4, 37]. These approaches cannot take into account any unforeseen events. Furthermore, preparing the routing for worst-case traffic demands may degrade performance on average [86]. Routing protocols can also take into account the result of traffic matrix estimations, on which there exists a variety of literature [64, 78, 86, 87].

Though off-line approaches offer good performance gains, they obviously cannot react to unforeseen short-term changes of traffic demands that may be caused by flash crowds, link failures, interdomain reroutes, etc. This necessitates the deployment of dynamic traffic engineering approaches. With MATE [25] and TeXCP [44] two systems were presented that work similar to our approach. Traffic is split among alternative paths and the splitting ratio is permanently being adjusted to the current traffic situation. For setting up path alternatives, the authors propose to use MPLS tunnels. Since every path has to be represented explicitly, this generates a communication overhead that does not scale well, since it is quadratic in the number of edge routers. Using implicit representations of paths and aggregate information about their performance, our algorithm can handle larger sets of alternatives more easily. Furthermore, these systems have not been evaluated under realistic conditions, including TCP feedback and bursty traffic.

Another approach similar to ours is the PSTARA protocol [59, 60, 75]. Here, the focus is on small mobile networks. Realistic traffic workload and TCP issues are not considered. Liu and Reddy analyse a similar adaptive mechanism in a multihoming scenario [51].

5.1.2 Notation

Throughout this chapter, we use the following notation. For any router node r and any destination t, the set $N(r,t)$ denotes the next-hop routers in the routing table of r admissible for destination t. Here, a destination t may denote a single node or an entire subnet. For the routing graph induced by the successor relation $N(\cdot,t)$ for any fixed destination t, we denote a (directed) edge from v to w by $v \to w$, and by $v \rightsquigarrow w$ we denote a (directed) path of arbitrary length between node v and node w. We assume that the routing graph is acyclic which is generally the case for routing protocols and properly configured routing tables.

As mentioned above, our protocol is generic with respect to the measure of performance used. For the sake of coherence with the preceding chapters, we will still use the symbol ℓ to denote the performance value with the convention that this value is to be minimised.

5.2 The REPLEX Protocol

Although the (α,β)-exploration-replication policy serves as an inspiration and eponym for our protocol, it cannot be translated into a traffic engineering protocol immediately, and we will deviate from it in a number of aspects. The Wardrop model which was used for our theoretic analysis made assumptions that are not met for the scenario of Internet routing. First, it assumes that the agents are capable of selecting an entire path for the packets they insert into the network. In networking terminology, this is referred to as *source routing*, which is available in specialised networks like certain overlay architectures, but not in pure IP networks. Second, what was phrased as "observing the latency of a path" in our theoretical descriptions of rerouting policies actually comprises a non-negligible effort of measuring and aggregating performance values in practise. Finally, the assumption of an infinite number of agents reduces to the assumption that traffic may be split arbitrarily at every node. Whereas this effect can be achieved by forwarding packets via probabilistically chosen next-hop routers, this may result in undesirable packet reorderings at the destination node. In the following sections, we will elaborate on the three above-mentioned problems and finally introduce our traffic engineering protocol formally.

5.2.1 Delegation of Routing Decisions

As opposed to the assumptions made in our theoretical model, we can now no longer assume that end users have full control over the path a packet takes on the way to its destination. Usually, end users are located at the periphery of the network and are connected via a single link only, e. g., using Ethernet or DSL. Consequently, end hosts do not have any influence on the routing. In regular IP routing not even the routers themselves can select a path as a whole. Unless using specialised protocols like MPLS tunnels, routers may forward packets only a single hop to bring them closer to their final destination. Thus, each router r narrows down the set of paths a packet can possibly take.

Consider a router r on a path connecting a source node s and a destination t. For every next hop router $v_i \in N(r,t)$, r maintains a weight $w(r,t,v_i)$, and we normalise $\sum_{v \in N(r,t)} w(r,t,v) = 1$. Given these weights, the router then aims at partitioning the traffic towards t such that the fraction going over $v_i \in N(r,t)$ is as close as possible to $w(r,t,v_i)$. The weights will be adapted over time depending on the performance of paths $r \to v_i \rightsquigarrow t$.

In order for router r to be able to assign reasonable weights on behalf of the end host s some performance information is required. Actually, the weight assigned to a path should depend on the performance of the entire path $s \rightsquigarrow r \to v_i \rightsquigarrow t$. However, once a packet has reached router r, the user's preference relation merely depends on the performance of the path section $r \to v_i \rightsquigarrow t$ and is independent of the section $s \rightsquigarrow r$. Since r cannot select a particular path $r \rightsquigarrow t$ as a whole, but can merely determine the next hop, r's decision should furthermore depend on aggregate performance information about the set of paths $r \to v_i \rightsquigarrow t$ only rather than on the performance of individual paths from this set.

5.2.2 Collaborative Performance Evaluation

How does the router r obtain this aggregate performance of path sections $r \rightarrow v_i \rightsquigarrow t$? Let us denote this aggregate performance by $L(r,t,v_i)$. The value of $L(r,t,v_i)$ is composed of two components: The performance of the edge $r \rightarrow v_i$, denoted ℓ_{rv_i}, and the overall performance of the set of paths $v_i \rightsquigarrow t$, denoted $A(v_i,t)$. The first component, ℓ_{rv_i}, can be measured directly at the router r. The second component, $A(v_i,t)$, however, must be measured and aggregated at v_i, and then transmitted to r. This may involve a recursive aggregation step at v_i. Finally, r may compute the overall performance $L(r,t,v_i)$.

How the aggregation is done exactly, depends on the type of performance measure we use. In general, we use a binary operator $\oplus$ to compute

$$L(r,t,v_i) = \ell_{rv_i} \oplus A(v_i,t)$$

and define $A(v_i,t)$ as the weighted average performance of alternative paths at v_i, i. e.,

$$A(v_i,t) = \sum_{u \in N(v_i,t)} w(v_i,t,u) \cdot L(v_i,t,u) \enspace .$$

Thus, the value of $L(\cdot)$ can be computed directly for routers in distance one from the destination and can then be propagated and aggregated backwards along the path.

For the case of using latency as our performance measure, we can use regular addition for the $\oplus$ operator since latency is additive along a path. It is intuitive and easy to check by induction on the path length that in this case, $L(r,t,v_i)$ is actually the weighted average latency (or the expected latency of a random packet) of all paths $r \rightarrow v_i \rightsquigarrow t$ where the weight of a path is the product of the weights of its edges. Let us remark that for the case that $x \oplus y = \max\{x,y\}$, $L(\cdot)$ does *not* yield the expected maximum of the performance values along a path. In fact, it is not possible to compute this value easily without increasing the amount of information transmitted between the routers. Still, we will see that even in this case traffic can be balanced based on the value of $L(\cdot)$.

What is the communication overhead incurred by our algorithm? Each router r transmits the value of all $A(r,t)$ for all destinations t in its routing table to its neighbours at regular intervals. Each entry of this message consists of an identifier of t, typically a IPv4 prefix ($32+5$ bits) or a IPv6 prefix ($128+7$ bits), as well as the value of $A(r,t)$. If we choose a resolution of 11 bits for this value, the total length of an entry is 6 bytes for IPv4 and less than 19 bytes for IPv6. Even if information about as many as 200 000 IPv6 prefixes is exchanged once in every two seconds, the communication overhead is only 1.8 MBytes/s which is less than 0.15% of today's link capacities of 10 Gbit/s. This can be further reduced by grouping subnet prefixes that are reachable via the same edge router.

Let us remark that though communication between the routers may be beneficial (see Section 5.4.4), the protocol can also be implemented without communication if this is not possible or desired.

5.2.3 Measuring Performance

The protocol described in the previous section is largely independent of the actual measure of performance used. We can measure or monitor various parameters with acceptable computational effort: the number of packets sent and dropped, the number of bytes sent and dropped, and the average queue length. From this information we can compute several performance measures:

1. The *latency* of an edge can be computed as the static propagation delay of a link plus the variable queueing delay. Along a path, latency is additive.

2. The *bit rate* of an edge is the number of bits sent divided by the length of the measurement period. The bit rate of a path is the minimum bit rate of its edges.

3. The *link utilisation* can be computed similarly. Let b denote the bit rate of an interface, let T denote the length of the measurement interval, and let n denote the number of bytes sent within this interval. Then, the link utilisation is $8\,n/(T\,b)$. We define the utilisation of a path to be the maximum utilisation of a link in the path.

Which of these parameters we choose to optimise depends on the type of application we want to support. Voice over IP requires low latency, accessing large files requires high throughput, and a typical goal of an Internet service provider is to minimise the maximum utilisation of a link within its AS. For our simulations we use link utilisation as our measure of performance. Consequently, we assume that the optimisation goal is to *minimise* the performance parameter.

Measuring the above values during short time periods will not yield satisfactory results since Internet traffic is known to be very bursty [57]. To this end, instead of using the measured values directly, we use an exponential moving average (EMA). More precisely, every router r maintains a variable $\tilde{\ell}(r,v)$ for every neighbour v, measures its current value ℓ_{rv} over a certain time period and updates

$$\tilde{\ell}(r,v) \leftarrow \eta \cdot \ell_{rv} + (1-\eta) \cdot \tilde{\ell}(r,v) \ ,$$

where $\eta \in (0,1)$ is a weighting parameter.

5.2.4 Randomisation vs. Hashing

Suppose we are given weights $w(r,t,v)$ for any destination t and any next-hop node $v \in N(r,t)$. How can we distribute traffic evenly to match these weights? A simple and natural answer to this question is to simply forward packets in a probabilistic fashion where the probability to choose a particular outgoing edge equals this edge's weight. However, in probabilistic routing it may happen that packets bound for the same destination overtake each other causing packet reorderings at the destination node. The TCP protocol treats such reorderings in the same way it treats packet losses, by decreasing the injection rate. Hence, the probabilistic approach may seriously harm performance [50].

This process can be derandomised by using a standard hashing technique [18]. Let V denote the set of identifiers of connection endpoints and consider a hash function $h : V \times V \mapsto [0,1]$ mapping each packet based on its source and destination to some value in the interval $[0,1]$. For a fixed destination t, we partition the interval $[0,1]$ into segments of size $w(r,t,v_i)$, $v_i \in N(r,t)$ and label each interval with the corresponding node v_i. A packet from source s for destination t is forwarded via the next-hop node associated with the interval containing $h(s,t)$. Note that each connection is not only characterised by the IP addresses of source and destination, but also by the corresponding port numbers. Thus, diversity can be enhanced significantly by using a combination of IP address and port number as the argument of the hash function.

Using this technique, we can ensure that no packet reordering occurs as long as the weights (and thus the intervals) do not change. Weight shifts cause a fraction of the traffic to be rerouted, and thus packets may be dropped. Hence, the time interval at which weight shifts occur should be larger than the amount of time it takes for TCP to recover from packet losses.

Whereas hashing is a good approximation of the probabilistic approach if the number of connections is large, hash values may be distributed unevenly if the number of connections is small. This may cause a next-hop node to receive more packets than intended. However, we get a solution to this problem for free since our algorithm will react to this effect by reducing the weight assigned to this unintentionally overloaded node. The effect that the number of end hosts has on the quality of the hashing will be examined in Section 5.4.3.

5.2.5 Adapting Weights

The actual adaption of the weights $w(r,t,v)$ depending on the performance values $L(r,t,v)$ resembles the (α,β)-exploration-replication policy (Equation (4.2)). However, in order to avoid weight shifts due to performance differences that are not significant, we introduce an additional threshold parameter ϵ. Fix a router r and a destination t. For any pair of next-hop routers v_1 and v_2, we shift weight from $w(r,t,v_1)$ to $w(r,t,v_2)$ if $L(r,t,v_1) > L(r,t,v_2) + \epsilon$. Corresponding to Equation (4.2), where we replace flow by weight, the amount of weight shifted from node v_1 to node v_2 is given by

$$\lambda \cdot w(r,t,v_1) \cdot \left((1-\beta)\cdot w(r,t,v_2) + \frac{\beta}{|N(r,t)|}\right) \cdot \frac{L(r,t,v_1) - L(r,t,v_2)}{L(r,t,v_1)+\alpha} \quad .$$

The migration speed can be controlled by the parameter λ and the ratio between exploration and replication is controlled by β.

5.2.6 Summary of the Protocol and its Parameters

Algorithm 5.1 Main loop of the REPLEX-algorithm

for every neighbour $v \in \bigcup_t N(r,t)$ **do**
  measure performance ℓ_{rv}
  // calculate exponential moving average with weight η:
  $\tilde{\ell}(r,v) \leftarrow \eta \cdot \ell_{rv} + (1-\eta)\tilde{\ell}(r,v)$
end for
initialise an empty message M
for every destination t in the routing table **do**
  for every next-hop node $v \in N(r,t)$ **do**
    $L(r,t,v) \leftarrow \tilde{\ell}_{rv} \oplus A(v,t)$.
  end for
  $avg \leftarrow \sum_{v\in N(r,t)} w(r,t,v)\cdot L(r,t,v)$
  Append (t, avg) to M
end for
send M to all neighbours v; these store it in $A(r,\cdot)$
call procedure ADAPTWEIGHTS (Algorithm 5.2)

Combining what has been presented in the preceding sections, the overall algorithm is summarised in Algorithms 5.1 and 5.2. Every router r maintains four arrays:

Algorithm 5.2 Procedure ADAPTWEIGHTS

for every destination t in the routing table **do**
  $w'(r,t) \leftarrow w(r,t)$.
  for every pair of next-hop routers $v_1, v_2 \in N(r,t)$ **do**
    if $L(r,t,v_1) > L(r,t,v_2) + \epsilon$ **then**
      $\delta \leftarrow \lambda \cdot w(r,t,v_1) \cdot \left((1-\beta) \cdot w(r,t,v_2) + \frac{\beta}{|N(r,t)|}\right) \cdot \frac{L(r,t,v_1) - L(r,t,v_2)}{L(r,t,v_1)+\alpha}$
      $w'(r,t,v_1) \leftarrow w'(r,t,v_1) - \delta$
      $w'(r,t,v_2) \leftarrow w'(r,t,v_2) + \delta$
    end if
  end for
  set $w(r,t) \leftarrow w'(r,t)$.
end for

$\tilde{\ell}(r,v)$**:** The exponential moving average of the measurements for link (r,v) for all neighbours of r.

$L(r,t,v)$**:** The performance value for destination t given that the next hop node is $v \in N(r,t)$.

$A(v,t)$**:** The average measurement value that next hop router v has announced for destination t. This array is updated whenever update messages from neighbouring routers are received. It is initialised with values that are neutral w. r. t. $\oplus$, e. g., with zero for latency or link utilisation.

$w(r,t,v)$**:** Current weight of route $r \to v \rightsquigarrow t$ for destination t. Initially, we choose uniform weights $w(r,t,v) = 1/|N(r,t)|$ for all $v \in N(r,t)$.

Observe that the memory requirement for this data is linear in the size of the routing table.

The router's main loop which is executed once in T seconds is given in Algorithm 5.1. In this loop, the performance ℓ_{rv} is measured for all neighbours v and the entries of $L(r,\cdot,\cdot)$ are updated accordingly. Then, for all destinations the weighted average performance over all next hop routers is computed and transmitted to adjacent routers which in turn update their table $A(r,\cdot)$. Finally, as is specified by Algorithm 5.2, all weights are updated.

Our algorithm can be parametrised in a number of ways. We summarise these parameters here and describe their influence on the behaviour of the protocol.

update period length T**:** The length of the interval at which the main loop is executed. This has a direct influence on how much traffic is rerouted within a time interval of a given length.

damping factor λ**:** This parameter determines the amount of weight shifted in one round. The ratio λ / T controls the convergence speed of the algorithm. Possible overshooting and oscillation effects limit the maximum convergence speed we can achieve.

EMA weight η**:** Decreasing the weight of the exponential moving average makes the algorithm less sensitive to the effects of bursty traffic whereas choosing it too small increases the time until the algorithm realises the effects of its rerouting decisions.

virtual performance offset α**:** This virtual offset is added to all path latencies $L(r, t, v)$ making the algorithm less sensitive to small differences when the performance value is close to 0. (For a detailed discussion, see Sections 4.1.3.2 and 4.2.) If α is too small, the algorithm shifts too much weight to paths that had not been used so far (e. g. when a link was down and is brought up again) which can cause oscillation. Choosing α too large makes performance differences less significant and reduces the speed of convergence. Note that the dimension of α is the dimension of the performance parameter ℓ.

improvement threshold ϵ**:** The parameter ϵ is used to avoid weight shifts between routes with insignificant performance difference. If this difference is below ϵ, no weight shift is performed. Thus, we can ensure that weights are only shifted if the change is substantial. Choosing ϵ too large limits the approximation quality we can achieve. Choosing it too small, the algorithm may tend to exhibit trembling behaviour. As with α, the dimension of ϵ is the dimension of the performance parameter ℓ.

exploration ratio β**:** This parameter determines the ratio between exploration and replication. If β is too small, currently unused paths may not be detected by our algorithm. On the other hand, β should not be too large since this can result in excessive growth of a flow if it is close to zero and no reliable performance information is available yet. (For a detailed discussion, see Section 4.2.)

Among the six parameters listed here, only three have substantial influence on the convergence behaviour of our protocol. The parameters α and ϵ are very similar. We have found that $\epsilon = \alpha = 0.1$ is a reasonable choice, but other values in the same order of magnitude do not change our results significantly. The role of the parameter β is less important in practise than it is in our theoretical analysis since, typically, weights do not become extremely small. In the following, we set $\beta = 0.1$.

The parameters λ, T, and η are more interesting since they determine the speed of convergence and the strength of oscillation effects. We explore the parameter space systematically in Section 5.4.

5.3 Simulation Setup

Our protocol was implemented and simulated within the SSFNet simulation framework [61]. In SSFNet, network traffic is simulated at packet level with a complete TCP implementation. Furthermore, it features several routing protocols including OSPF, MPLS, and BGP.

The purpose of our simulations is twofold. On the one hand, we want to find a good set of parameters as described in the previous section. To that end, we have used simple scenarios consisting of a few routers only. These scenarios were designed to observe specific effects. On the other hand, we want to understand the behaviour of our protocol in complex real-world scenarios. To that end, we have used complex AS topologies provided by the Rocketfuel [3, 76] dataset. The specific topologies are discussed within the respective sections. The workload used is the same for all topologies and is described below.

All topologies consist of a set of routers, some of which are connected to so-called *workload clouds* consisting of Web clients and Web servers that generate the actual traffic demands.

5.3.1 Realistic Web Workload and Traffic Demands

One of the key aspects in which our simulations differ from earlier results is the fact that we use a realistic workload generator taking into account the bursty nature of traffic on the Internet [57]. We believe that standard approaches like replaying a pre-recorded trace do not suffice, as they do not account for interacting feedback loops: Our traffic engineering protocol and other control loops, foremost TCP, interact with each other and may influence the traffic demands themselves. It is crucial to take feedback loops of this kind into consideration for the evaluation.

To generate realistic HTTP traffic, we use the Web workload generators of [81] and [61] which simulate the behaviour of HTTP clients. The statistical properties of the workload model are based on measurements of real Internet traffic [30] and are similar to the one imposed by SURGE [10] or NSWeb [82]. To increase burstiness, we disable the use of persistent HTTP connections [48].

Both generators simulate Web users that request a Web page, then are idle for some time before reading another Web page, etc. Both the timeouts and the simulated file sizes are heavy-tailed and thus yield many short flows, as well as a significant number of long-lasting flows. The Web pages chosen by the users for download are chosen uniformly at random from any Web server. Hence, the workload demands are uniformly distributed and reflect the traffic matrix model of Zhang *et al.* [87].

5.3.2 Initialisation

Each simulation starts with an initialisation phase of 500 seconds. This phase is used to wait for the underlying routing protocol to converge (if such a protocol is used, as is done in Section 5.5). This happens within seconds. After the routing is established, the workload generating clients described above wake up at random points of time to avoid synchronisation effects. After the initialisation phase, the evaluation phase starts. In all figures presented below, the time axis starts at the beginning of the evaluation phase.

5.4 Exploring the Parameter Space using Artificial Topologies

In this Section, we use artificial topologies consisting of a few routers only. No underlying routing protocol is used and routing paths are set up manually as will be described in the respective sections.

5.4.1 Convergence Speed vs. Oscillation

Both our theoretical analyses in the previous chapters and practical experience indicate that there is a tradeoff between the speed of convergence we can achieve and the strength of oscillation effects. Convergence can be accelerated by increasing the frequency $1/T$ at which weights are adapted and by increasing the parameter λ which determines the maximum amount of flow shifted per round. Increased convergence speed comes at the cost of an increased liability to oscillation. The role of the weight parameter η of the exponential moving average plays an ambivalent role in this context. If we chose η too large, our protocol will overreact to accidental traffic fluctuations. If we chose it too large, a systematic change of performance will take too long to have an effect. Hence, both extremes may lead to oscillation.

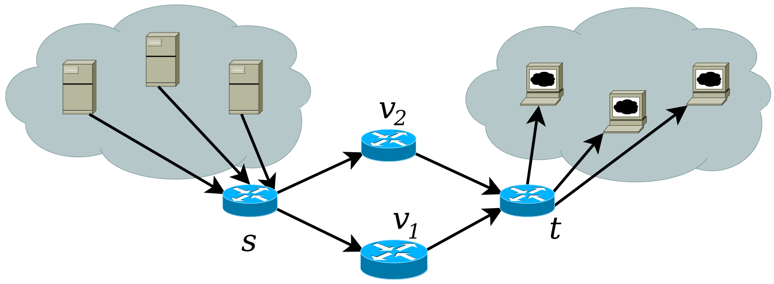

Figure 5.1: A set of HTTP servers is connected to a set of HTTP clients via two routes $s \to v_1 \to t$ and $s \to v_2 \to t$ with bandwidths S_1 and S_2, respectively.

We chose a simple topology depicted in Figure 5.1 to find suitable parameters T, λ, and η. The network consists of four routers s, v_1, v_2, and t, where s acts as a source being connected to 40 HTTP servers and t acts as a sink being connected to 40 HTTP clients. From s, t can be reached via two paths $s \to v_1 \to t$ and $s \to v_2 \to t$. Both links of one path have equal bandwidth, and we denote the bandwidths of the paths by S_1 and S_2, respectively. Consequently, in order to balance link utilisation, the traffic should be split according to the ratios $S_1/(S_1 + S_2)$ and $S_2/(S_1 + S_2)$.

For our evaluation we use unbalanced bandwidths $S_1 = 10\,\mathrm{Mbit/s}$ and $S_2 = 100\,\mathrm{Mbit/s}$. For these bandwidths, the optimal weights that balance the link utilisation are $w(s,t,v_1) \approx 0.09$ and $w(s,t,v_2) \approx 0.91$. We have configured the propagation delays to be 10 ms, but the results presented here are largely independent of this quantity. Experiments were also conducted with propagation delays ranging from 0.1 ms to 100 ms, but the behaviour of our protocol was unaffected by this.

Figure 5.2 shows the results of our simulations for various values of λ, T, and η. We plot the weights of both routes over an interval of 1,500 s. Note that $w(s,t,v_1) + w(s,t,v_2) = 1$, so the information given in the plots is partially redundant. As expected, the critical parameter is $\tilde{\lambda} = \lambda/T$. As $\tilde{\lambda}$ increases, the speed of convergence improves at the cost of stronger oscillation effects.

In the first plot of Figure 5.2 we have $\lambda = 0.05$ and $T = 8s$, so $\tilde{\lambda} = 0.00625\,\mathrm{s}^{-1}$. We see that convergence is very smooth for all values of η, but takes 750 s to approach its final value. Increasing $\tilde{\lambda}$ by a factor of 6 to $0.4\,\mathrm{s}^{-1}$ ($T = 1\,\mathrm{s}$, $\lambda = 0.4$), the last plot shows the other extreme. Here, the weight of the first route oscillates rapidly between 0 and 0.4 for all values of η. However, even in this extreme case, the time average of the weights is still close to the optimal values.

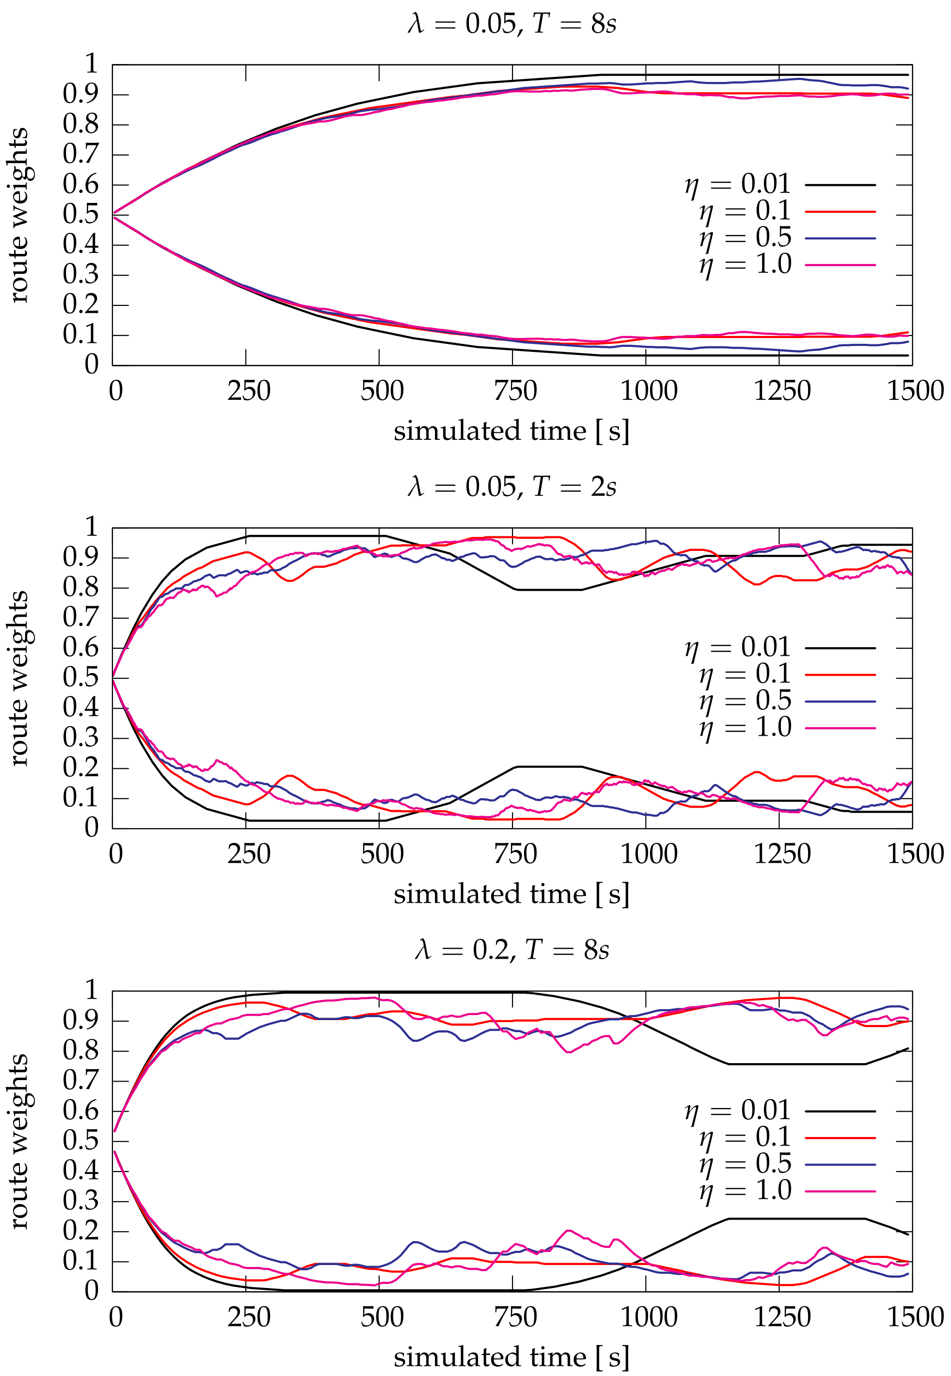

Figure 5.2: The above plots show how weights are adapted for various values of λ, T, and η over a time period of 1,500 s in the network depicted in Figure 5.1 with $S_1 = 10$ Mbit/s and $S_2 = 100$ Mbit/s. For every choice of parameters, we have two plots of weights, one for each route. As the ratio $\tilde{\lambda} = \lambda/T$ increases, we observe a speed-up in convergence at the cost of stronger oscillation.

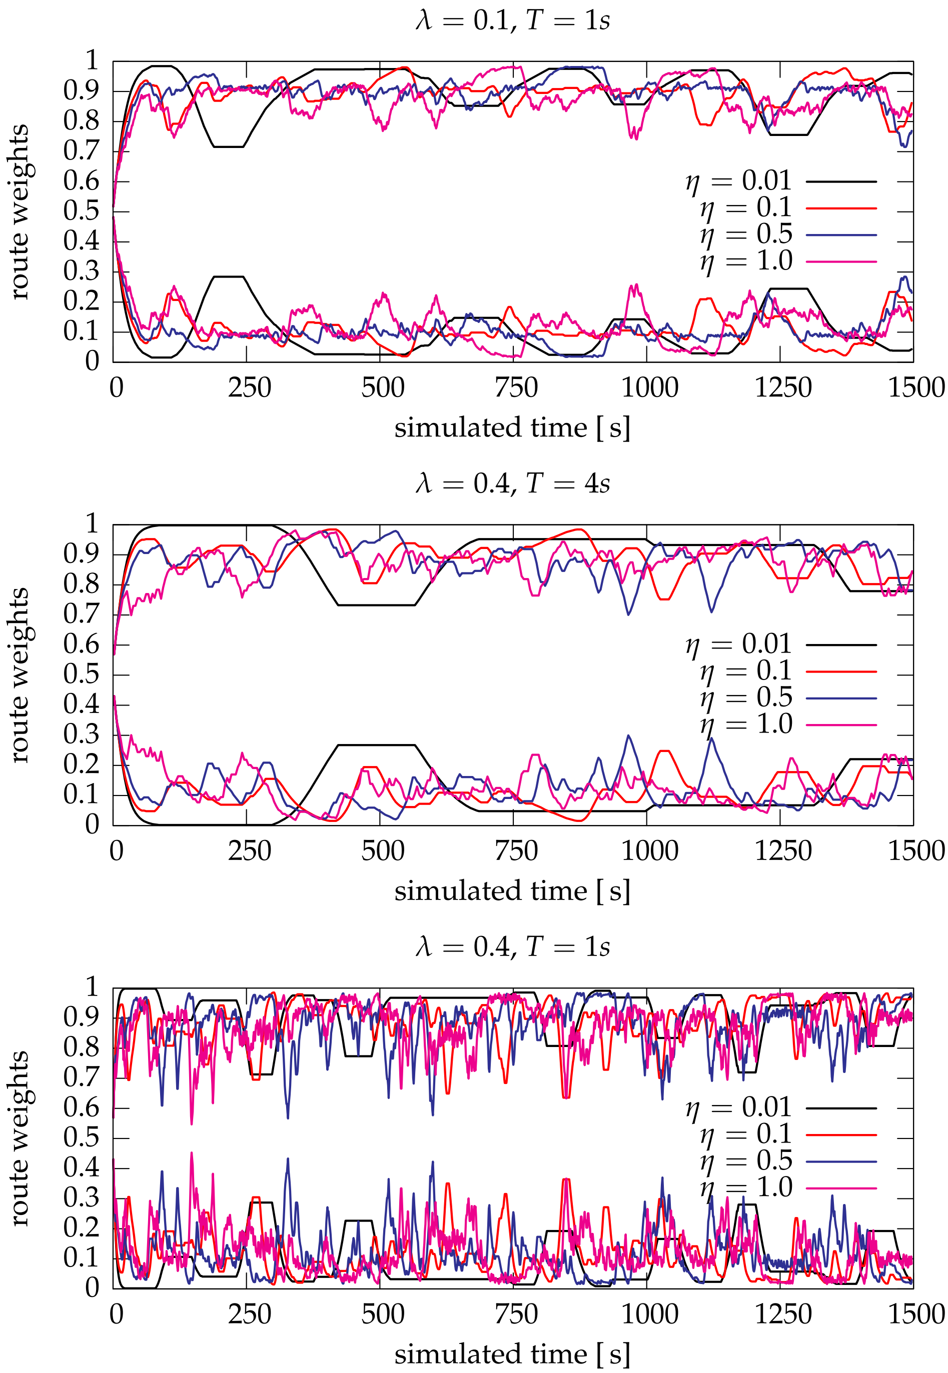

Figure 5.2: (Continued)

We observe good behaviour of the system for the case that $\tilde{\lambda} = 0.025\,\text{s}^{-1}$. In Figure 5.2 we have depicted the cases $T = 2\,\text{s}$, $\lambda = 0.05$ and $T = 8\,\text{s}$, $\lambda = 0.2$. For these parameter values, we have relatively fast convergence while at the same time keeping oscillation at an acceptable level. Considering the two sets of parameters with $\tilde{\lambda} = 0.025\,\text{s}^{-1}$ we can observe that for longer measurement intervals T and larger values of λ, the amplitude of the oscillation increases slightly. Hence, we should choose T as the minimum amount of time such that measurements are still reasonable and then choose λ accordingly to reduce oscillation. Note, however, that reducing T increases the communication overhead of our protocol. Since this overhead is negligible (see Section 5.2.2), the improvement in oscillation avoidance weighs out the additional communication overhead.

Our simulations also reveal the impact of the weight η of the exponential moving average. Small values of η smooth out peaks and thus stabilise our protocol. However, if η is chosen too small, the protocol becomes inert since it takes a longer time to realise significant performance changes.

Summarising our observations, we find that $T = 2\,\text{s}$, $\lambda = 0.05$, and $\eta \in (0.1, 0.5)$ is a good parameter set. These parameter settings achieve convergence within 200 s and reduce oscillation to an acceptable amount. We use these parameters ($T = 2\,\text{s}$, $\lambda = 0.05$, $\eta = 0.1$) in the simulations presented below.

5.4.2 Quality of the Solution

In the preceding section we have used alternative routes with a rather large discrepancy between the bandwidths of the links. Using the parameters determined above, we now use the simple topology depicted in Figure 5.1 again, but vary the bandwidths of the links. We fix $S_1 = 100\,\text{Mbit/s}$ and vary S_2 to be 10 Mbit/s, 34 Mbit/s, and 100 Mbit/s. Consequently, in order to balance link utilisation, optimal weights for both routes are approximately 0.91/0.09, 0.75/0.25, and 0.5/0.5. Figure 5.3 shows that our algorithm actually converges towards these sets of weights (upper plot), and that link utilisation becomes actually balanced (lower plot). The total flow demand was chosen to be approximately 10 % of the total bandwidth available. Due to our careful choice of parameters, oscillation can be avoided.

5.4.3 Necessary Number of End Hosts

In Section 5.2.4 we have pointed out that our hashing technique for derandomisation is a good approximation of probabilistic routing only if the input range of the hash function is large, i. e., if there exist a sufficient number of connections. How many connections do we need?

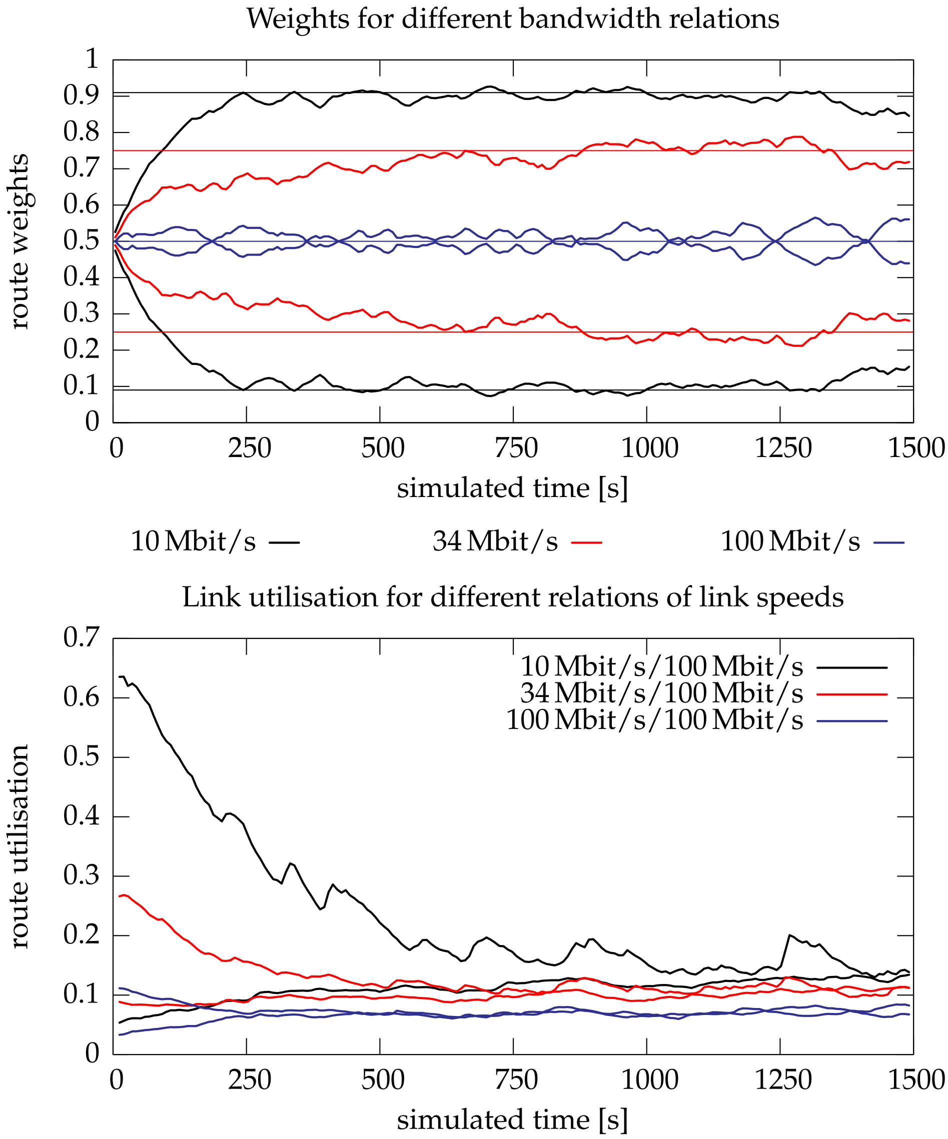

Figure 5.3: The plots show the weight and the link utilisation in the network depicted in Figure 5.1 over an interval of 1,500 s. The bandwidth of one path was fixed to $S_1 = 100$ Mbit/s while the other was set to $S_2 = 10$ Mbit/s, 34 Mbit/s, and 100 Mbit/s. We see that in all cases the link utilisation becomes balanced over time. Thin lines indicate the optimal weights.

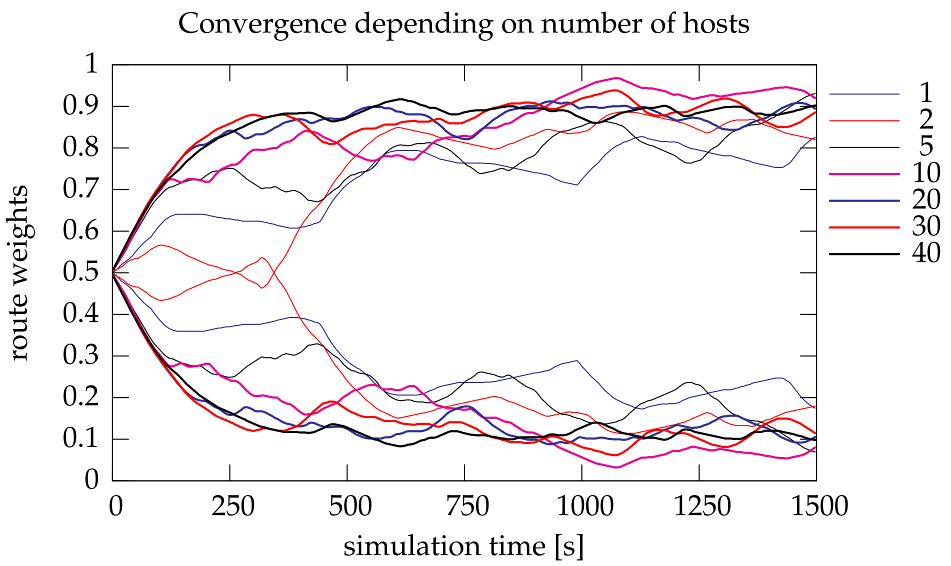

Figure 5.4: The protocols ability to find optimal weights depends on the number of end hosts n, since these determine the number of distinct connections that are inputs to the hash function.

To answer this question, we use the simple topology depicted in Figure 5.1 again, setting the bandwidths to $S_1 = 10\,\text{Mbit/s}$ and $S_2 = 100\,\text{Mbit/s}$. We vary the number n of servers and clients connected to the routers s and t. Note that for n end hosts there exists a total number of $\Theta(n^2)$ possible connections.

Using the parameter set determined above, Figure 5.4 shows how our algorithm adapts weights over time. We see that for $n \in \{1, 2, 5\}$, the algorithm does not exhibit good convergence behaviour. Increasing the number of end hosts to $n \geq 10$, however, we see that the optimal weights 0.09 and 0.91 are approached and oscillation vanishes as n is increased.

5.4.4 Communication Helps

In the simple topology we have used so far, the utilisation of the two links $s \rightarrow v_1$ and $v_1 \rightarrow t$ of the first path is approximately equal. The same holds for the links of the second path. Hence, any information available at v_1 and v_2 is also available at s, implying that s cannot gain any additional benefit from performance messages transmitted by the neighbouring routers. To examine a topology in which communication between the routers is actually necessary, we consider the topology depicted in Figure 5.5. The router s under study is now connected to two destination routers t_1 and t_2 each of which is connected to a different set of 40 end hosts. Both destinations can be reached via intermediate routers v_1 and v_2. The bandwidth of links adjacent to s ($s \to v_1$ and $s \to v_2$) is $S_1 = 100\,\text{Mbit/s}$. The destination routers are connected to their respective intermediate routers via 100 Mbit/s links, too ($v_1 \to t_1$ and $v_2 \to t_2$). The cross links $v_1 \to t_2$ and $v_2 \to t_1$ are bottleneck links with bandwidth $S_2 \in \{10\,\text{Mbit/s}, 34\,\text{Mbit/s}, 100\,\text{Mbit/s}\}$. From the perspective of s, the bottleneck links are one hop away, and for t_1 and t_2 bottlenecks are behind different first hops. Hence, s needs to rely on performance information transmitted by v_1 and v_2. Again, optimal weights are $S_1/(S_1 + S_2)$ and $S_2/(S_1 + S_2)$, but for the different destinations t_1 and t_2, the intermediate routers v_1 and v_2 take opposing roles.

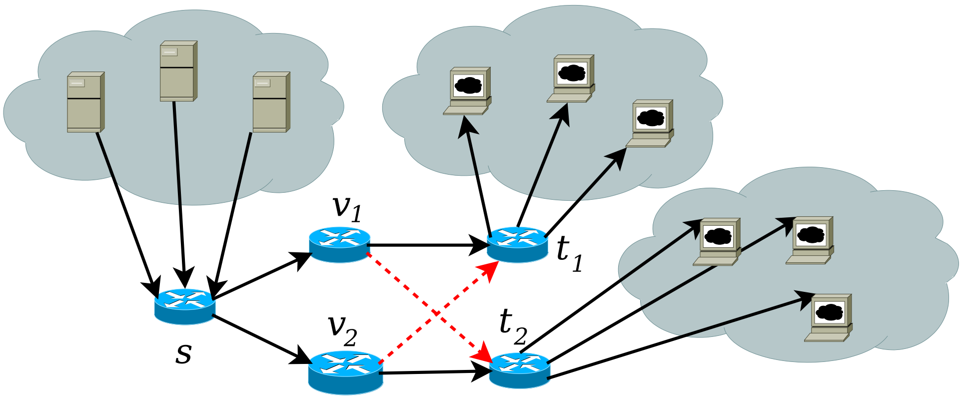

Figure 5.5: Servers at router s are connected to two sets of clients connected to t_1 and t_2. Each target router is reachable via two alternative routes. Bottleneck edges are red and dashed.

In this setting, the router s has to maintain four weights since there are two routes for each destination ($w(s, t_1, v_1)$, $w(s, t_1, v_2)$, $w(s, t_2, v_1)$, and $w(s, t_2, v_2)$). Figure 5.6 shows how these weights are adapted over time. For every value of S_2, we see that our algorithm approaches the optimal weights $S_1/(S_1 + S_2)$ and $S_2/(S_1 + S_2)$ quickly. Let us remark that, if communication between the routers is deactivated, our protocol cannot find good weights. In that case, all weights remain close to 0.5 throughout the simulation.

5.4.5 Drifting Demands

Though in the scenarios simulated so far, the traffic actually generated was bursty, the traffic demands themselves were static. Whereas static traffic demands can be optimised for using off-line algorithms, a traffic engineer-

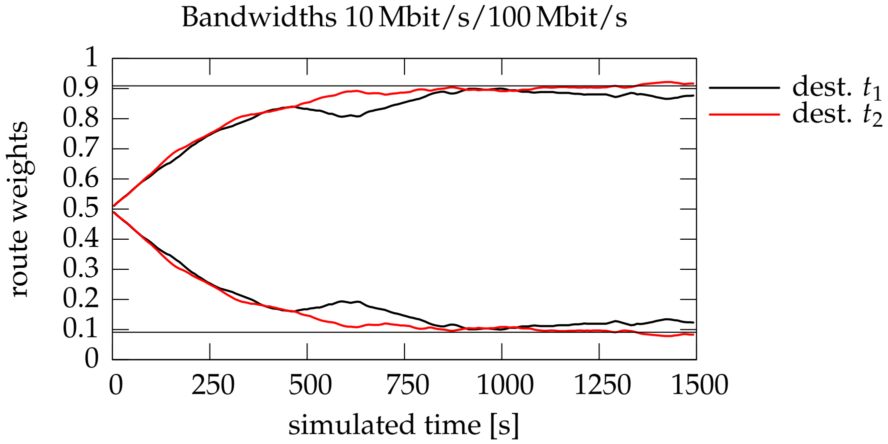

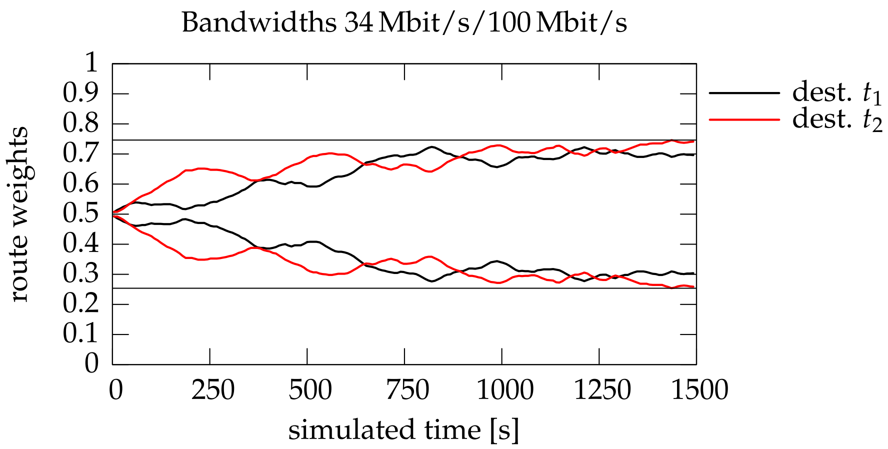

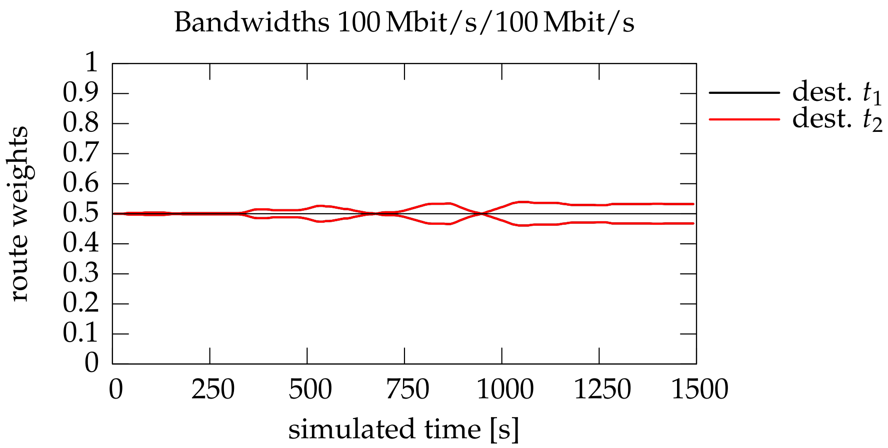

Figure 5.6: Adaption of weights for two destinations with bottlenecks that are not directly observable (Figure 5.5).

ing protocol should be able to react to changing traffic demands on-line. To simulate such a scenario, we use the topology depicted in Figure 5.7. Again, router s can connect to destination t via paths $s \to v_1 \to t$ and $s \to v_2 \to t$. All links have a bandwidth of 100 Mbit/s. This topology differs from the simple one used earlier in that there is a second set of clients connected to the intermediate router v_2. During a startup phase of $T_1 = 500$ s, these clients are inactive. Only at time T_1 do they start to wake up and also connect to the servers via s. The wake-up time of a client is chosen uniformly at random from the interval $[T_1, T_2]$ where the parameter $T_2 > T_1$ is varied. This increases the load on link $s \to v_2$ which should in turn lead to a decrease of the corresponding weight $w(s, t, v_2)$. Note, that s does not have routing alternatives for destination v_2.

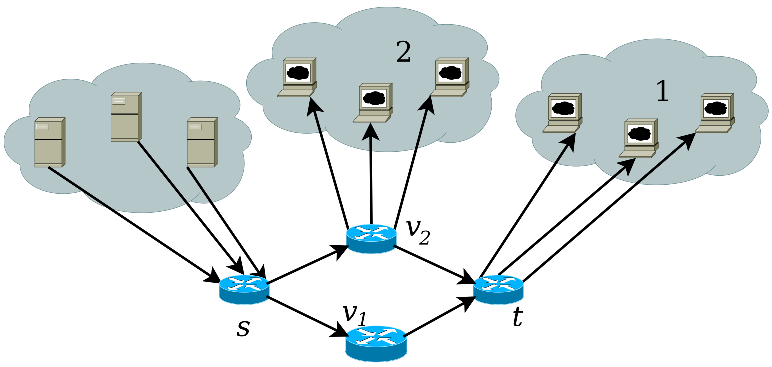

Figure 5.7: A set of D servers is connected to router s and can reach the target router t connected to the clients via intermediate routers v_1 and v_2. At time $T_1 = 500$ s, a second set of D/k servers connected to router v_2 starts to wake up.

Denote the number of hosts connected to t by n, and denote the number of hosts connected to v_2 by m. Furthermore, assume that the traffic demand is proportional to the number of clients. Let $w = w(s, v_1, t)$. Then, the traffic demand on link $s \to v_1$ is proportional to $n \cdot w$, and the traffic demand on link $s \to v_2$ is proportional to $n \cdot (1 - w) + m$. Consequently, the flow is balanced if

$$n \cdot (1 - w) + m = n \cdot w \ ,$$

or, equivalently, if $w = \min\{(m + n)/(2n), 1\}$.

Figure 5.8 shows the results of the simulation for $n = 40$ and $m = 10, 20$, and 40, i. e., the optimal value of w after all clients have woken up is 0.62, 0.75, and 1, respectively. The lengths of the wake-up periods were chosen to be $T_2 - T_1 = 10$ s, 100 s, and 500 s. We see that for long wake-up periods, e. g., $T_2 - T_1 = 500$ s, weights follow the changing demands smoothly. If demands change abruptly, the algorithm takes its regular convergence time to adapt to this change.

5.5 Simulation in a Real-World Topology

To evaluate the performance of the REPLEX protocol in a real-world network we chose the EBONE network from the Rocketfuel [3, 76] topology dataset (AS1755). This topology consists of 172 routers connected by 367 links. Out of the 172 routers, 61 are access and border routers. The topology is depicted partially in Figure 5.9. Due to its good connectivity, there exist various paths with equal hop count between pairs of routers, and our TE protocol can take advantage of this diversity. Altogether, we connect 5,400 clients and 417 servers to the routers. Our simulation setup is described in more detail in the following sections.

5.5.1 Network Topology and Client and Server Location

The Rocketfuel topologies contain only information about the routers of an AS. At every border and access router we install workload clouds consisting of a set of clients and servers. It is reasonable to assume that routers that are well-connected to the AS are also connected to many end hosts and thus are responsible for a large amount of traffic. Therefore, we chose the number of end hosts connected to a router proportional to its degree. More precisely, we install 30 clients and three servers for each link into the AS.

The Rocketfuel data set provides information about topology and estimated latencies of the links. Bandwidth information, however, is not provided. We use the propagation delays obtained from the Rocketfuel data set and configure the bandwidth of intra-AS links manually. In order to simulate an under-provisioned network which necessitates traffic engineering, we choose the bandwidth of such links to be small, namely 10 Mbit/s. End hosts are connected to the access routers by faster links with a bandwidth of 40 Mbit/s. Hence, bottlenecks are located within the AS and not at the periphery. The link delays of end hosts are randomised to add variability to the experienced round trip time [30]. They are chosen uniformly at random from the interval $[0, 60]$ ms for clients and $[0, 25]$ ms for servers.

5.5.2 Routing

Our traffic engineering protocol can operate on top of any routing protocol that can provide multiple paths for one destination. Whereas these paths could be configured manually for our simple artificial topologies, we use

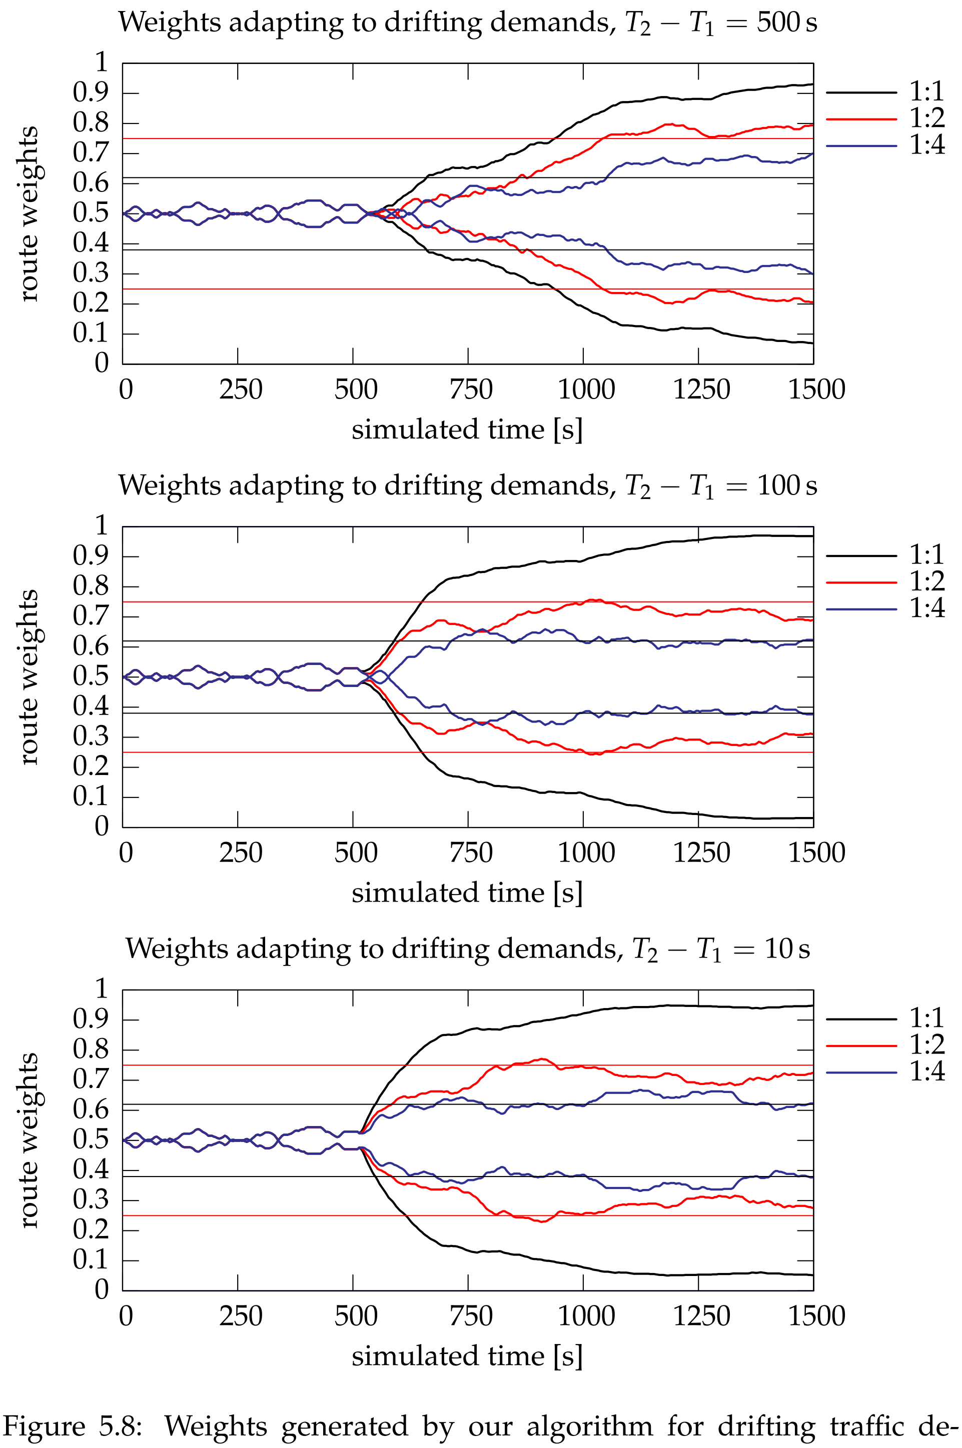

Figure 5.8: Weights generated by our algorithm for drifting traffic demands. A set of m clients at v_2 (see Figure 5.7) wake up between T_1 and T_2. The additional traffic demand is given as ratio $1 : 1$, $1 : 2$, and $1 : 4$ relative to the number of clients at t.

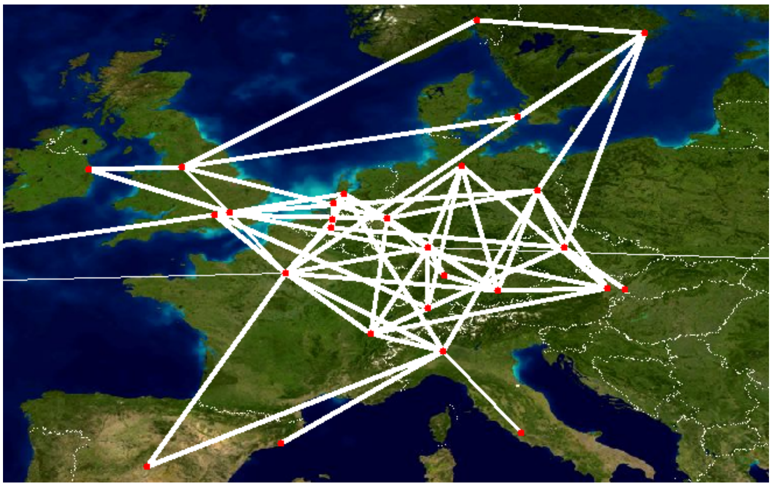

Figure 5.9: The Rocketfuel topology for EBONE (AS 1755, `.r0` map) consists of 172 routers. In total we connect 5,400 Web clients and 417 Web servers to the border and access routers.

the Open Shortest Path First (OSPF) protocol with uniform link weights for the Rocketfuel topology. Thus, hop count is our distance metrics, and all paths with equal distance from the destination form a set of admissible routes. This provides large flexibility for the traffic engineering protocol.

5.5.3 Performance Evaluation

We evaluate the performance of our protocol using the parameters $\lambda = 0.3$, $T = 2.5\,\mathrm{s}$, $\eta = 0.15$, and we perform simulations with communication between the routers enabled and disabled. We are interested in two properties: How fast does the algorithm converge, and how does it improve the throughput of the network?

To answer the first question, we consider the aggregate change of all weights $w(\cdot,\cdot,\cdot)$ at all routers. We sum up these changes over intervals of five seconds and normalise such that the largest value is 1. Figure 5.10 shows that this value decreases quickly over time. We see that significant weight changes occur only during the first two minutes, and the minimum is reached after three minutes. Note that due to the burstiness of the traffic weight fluctuations do not vanish completely. Observe that for the case of disabled communication, these fluctuations are greater than with communication enabled.

From the point of view of the end users, convergence of the routing protocol is not necessarily the primary issue. A traffic engineering pro-

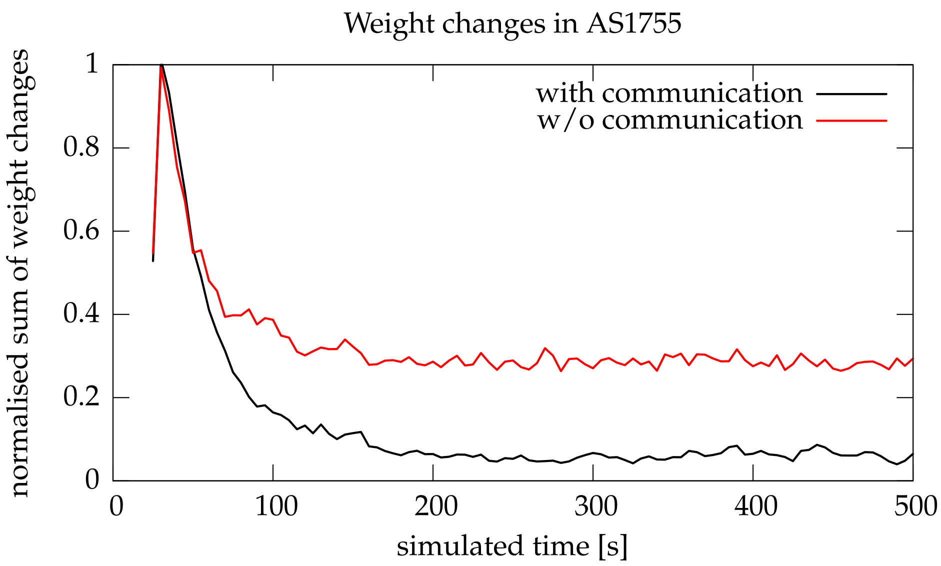

Figure 5.10: Aggregate weight changes in AS1755.

tocol presents a benefit to end users only if it improves the performance experienced by them. We measure this performance in terms of bandwidth per transferred Web object. We measure TCP throughput (in Byte/s) and consider the average and median over intervals of 20 seconds length. The results are depicted in Figure 5.11.

First observe that, due to the burstiness of Web traffic, mean and median throughput differ significantly. Whereas after our startup period of 500 s, median throughput is 23.5 kByte/s, the average throughput is almost 42 kByte/s. Even with communication disabled, median and mean throughput grow by 29 % to 30.5 kByte/s and by 14 % to 48 kByte/s. With communication enabled, the benefit is even larger. Median throughput increases by 59 % to 37.5 kByte/s and average throughput increases by 34 % to 56 kByte/s. The latter values have been calculated as time averages over the time interval [2000 s, 3500 s], i. e., over an interval at which the algorithm has converged. As is indicated by the decrease of weight shifts, a significant throughput gain is achieved already within the first three minutes, but the throughput keeps increasing slightly even afterwards.

5.6 Summary

We have introduced REPLEX, a dynamic traffic engineering protocol that balances network traffic on-line. Simulations have shown that our protocol is able to cope with the risk of oscillations, one of the major obstacles

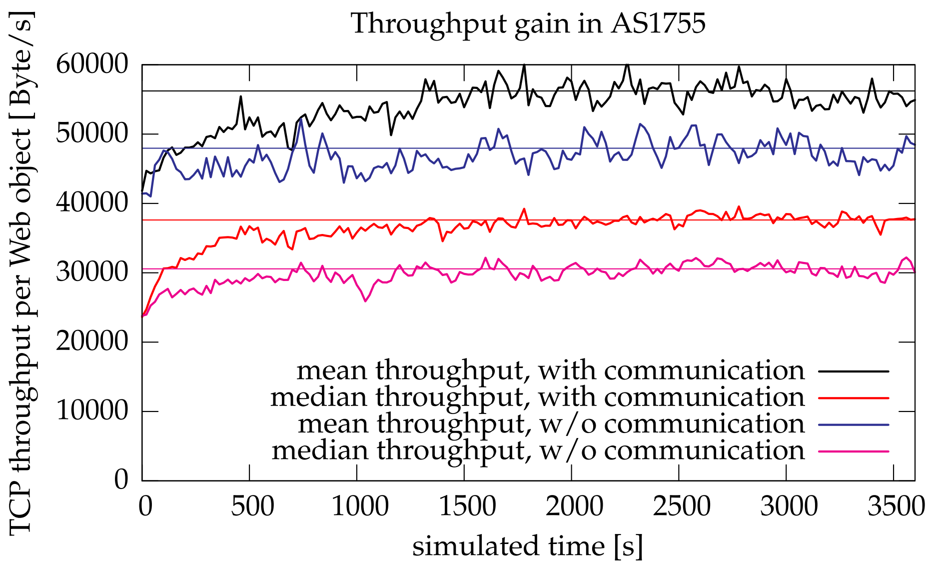

Figure 5.11: Throughput gain in AS1755.

that prevents dynamic traffic engineering protocols from being deployed in practise. Via simulations on artificial topologies we have determined a suitable parameter set which enables the protocol to converge within few minutes. On a real topology taken from the Rocketfuel data set with realistic workload, convergence comes with a significant improvement of the overall network performance measured in terms of throughput.

One of the advantages of our protocols is that it is easy to deploy on top of virtually any routing protocol, be it OSPF, MPLS, or even BGP, as long as this protocol can provide multipaths. The underlying routing structure does not have to be changed or entirely replaced as is necessary for many TE protocols. We have chosen OSPF since it is easily possible to obtain a diverse set of paths using this protocol. This permits also a compact representation of these paths and it is not necessary to store them explicitly, e. g., as MPLS tunnels. A second advantage of our protocol is that it can be easily adapted to the optimisation goal of the network operator by choosing an appropriate performance measure.

We consider the protocol as it is presented here rather a framework than a detailed implementation. This framework leaves room for improvements, such as the design of strategies to coordinate weight shifts in a way that minimise the number of reroutings that the individual connections are subject to. The fact that our protocol performs well already at this rather coarse level of implementation shows that dynamic traffic engineering is possible and is a good indicator that it is worth investigating protocols of this flavour in more detail.

CHAPTER 6

Conclusions and Open Problems

6.1 Conclusions

In this thesis, we have considered load-adaptive rerouting policies in various models which enable a population of selfish agents to approach Wardrop equilibria. What can we learn from the results of this work? There are two take home messages. The first is structural, and the second is technical:

Wardrop equilibria require only mild assumptions. The classical, static motivation of Nash equilibria in strategic games requires strong assumptions: complete and accurate knowledge of the payoff function, common knowledge of unbounded rationality, and unbounded reasoning capabilities. Our dynamical analyses show that these assumptions are not necessary for selfish routing games in the Wardrop model.

We have seen that a population of selfish agents can attain a Wardrop equilibrium in the long run by following simple selfish rules. This merely requires that the latency functions truly reflect the preferences of the agents and that agents have the capability to observe the latency and flow values induced by the population as well as to perform very simple sampling and comparison steps. However, this dynamic population model still assumes perfect information in that any action of the agents can be observed immediately.

Subsequently, we have seen that the assumption of perfect information can be abandoned as well. Observations may be made with a delay as long as the adaption policy is not too greedy. From our analysis it also follows that the flow values may even be observed with a certain error or uncertainty. E. g., as long as the flow values observed by the agents are within constant factors of their true value, our analyses of convergence and convergence time still hold up to constant factors. Our general framework of α-smooth adaption policies requires the migration rates to be bounded by the reciprocal of the maximum slope of the latency functions and the maximum age of the information. This slows down the convergence process significantly. In order for the population to approach an approximate Wardrop equilibrium we need to assume that network parameters and flow demands themselves change slowly, if at all. Thus, we have traded the assumption of perfect information for the assumption of inertness of the underlying network.

To show that not even this assumption is necessary, we have considered a specific adaption policy, the exploration-replication policy. We have seen that due to a clever sampling technique, this policy reaches approximate equilibria very quickly. In particular, the time to reach a state in which at most a small fraction of the agents deviates by more than a small fraction of the average latency from the average latency of their commodity is almost independent of the network size.

These results strengthen the concept of Wardrop equilibria showing that they can be motivated using very mild assumptions only.

Agent sampling is powerful. The main technique that allows for our polynomial bounds on the convergence time of the exploration-replication policy is the fact that the policy utilises flow proportional sampling. This has two advantages: First, it avoids the possibly exponentially large factor $\Omega(|\mathcal{P}|)$ which would occur for all static sampling probabilities. Second, it allows us to use the elasticity as the crucial parameter of our analysis. This is possible only since elasticity measures the relative growth of a function value due to a relative growth of the argument. The growth of a flow, however, is relative only if we apply flow proportional sampling.

6.2 Open Problems

The theoretical part of this work leaves open one major question. Whereas we have shown in Chapter 4 that $(1+\epsilon)$-approximations of Wardrop equilibria can be attained by the exploration-replication policy in polynomial time for symmetric instances, it is not clear whether this is possible for asymmetric instances. Technically, the argument which fails for the asymmetric case in our proof is the fact that the average latency cannot change significantly without decreasing the potential (Lemma 4.9). For instances with several commodities (and therefore also several averages), this argument clearly fails. However, there is also a structural reason which indicates that processes similar to the one considered here may take long to approximate Wardrop equilibria in this sense. For discrete asymmetric network congestion games it has been shown that computing a pure Nash equilibrium is PLS-complete [1, 29]. Furthermore, it is known that any sequence of selfish steps may be exponentially long. Therefore, it is unclear whether the continuous Wardrop model allows for a polynomial convergence time in the asymmetric case.

A second open question is whether our policy can be applied algorithmically. This is problematic for two reasons. First, our model involves an infinite number of agents, and second, our policy requires to sample from a possibly exponentially large set of paths. Thus, simulating the exploration-replication policy in a naive way would require to maintain an exponential number of variables $(f_P)_{P \in \mathcal{P}}$. A way to circumvent this has been exposed lately in [**?**]. Therein, we consider a distributed algorithm executed by a finite number of agents, one for each commodity. Each agent may split its flow demand and strives to balance the latency of paths of its own commodity. This algorithm represents path flows only implicitly, namely as a natural decomposition of an edge flow vector in a DAG. This way, flow updates can be computed in polynomial time. The algorithm proceeds by sampling a path with a probability distribution given by the implicit path flows. If this path has above average latency, a large fraction of its bottleneck flow is redistributed among the other paths. Thus, we can achieve convergence results towards bicriterial equilibria that are similar to the ones presented in Chapter 4.

Finally, a challenging open question is whether convergence results of a similar kind can be obtained in the case of discrete agents. For networks of parallel links some convergence results have been presented, e. g., in [12, 28], but these analyses are restricted to linear or even identical latency functions. Whereas our model takes into account concurrency, it neglects probabilistic effects. Whereas the fluid limit model allowed us to identify random variables (like the fraction of agents utilising a particular path) with their expectations, this is not possible if the number of agents is finite. Furthermore, the various discrete processes that have been used to motivate the replicator dynamics and its refinements behave differently if the number of agents is finite.

Appendix A

Omitted Proof

Proof of Fact 4.2. We compare L and Φ termwise. For the upper bound consider a resource e and note that the contribution to the average is $\ell_e(f_e) \cdot f_e$ whereas the contribution to the potential is $\int_0^{f_e} \ell_e(u)\, du \leq \ell_e(f_e) \cdot f_e$ by monotonicity of ℓ_e.

For the lower bound, first note that both the ratio L/Φ and the elasticity are invariant under scaling of any type. More precisely, if the function $f(x)$ has elasticity d and $f(1)/\int_0^1 f(u)\, du = \rho$, then the function $\tilde{f}(x) = a \cdot f(b \cdot x)$ for any $a, b > 0$, has elasticity d and $\tilde{f}(1/b)/\int_0^{1/b} \tilde{f}(u)\, du = \rho$, too. Consequently, we can simplify the analysis by considering only those functions with $f(1) = 1$. We construct a family of piecewise linear functions $(f_{d,n})$ for $n > d$, $n \in \mathbb{N}$ with $f_{d,n}(1) = 1$ that satisfies the following two properties:

1. Let ℓ be any function with elasticity d and $\ell(1) = 1$. Then $\ell(x) \geq f_{d,n}(x)$ for any $x \in [0, 1]$.

2. $\int_0^1 f_{d,n}(u)\, du \geq 1/(d+1) \cdot (1 - o(1))$ where the $o(1)$ is with respect to n.

Since the two claims hold for any n, they imply the lower bound also on ℓ, i. e., $\int_0^1 \ell(u)\, du \geq 1/(d+1)$ which is our claim. For $n > d$, we define the piecewise linear function $f_{d,n}$ by specifying the value $f_{d,n}(x)$ only for the points $x_i = i/n$, $i = 0, \ldots, n$:

$$f_{d,n}(x_i) = \begin{cases} 1 & \text{if } i = n \\ f_{d,n}(x_{i+1}) \cdot \left(1 - \frac{d}{i+1}\right) & \text{if } d \leq i < n \\ 0 & \text{if } i < d \ . \end{cases}$$

For an illustration of $f_{d,n}$ see Figure A.1. The construction implies that

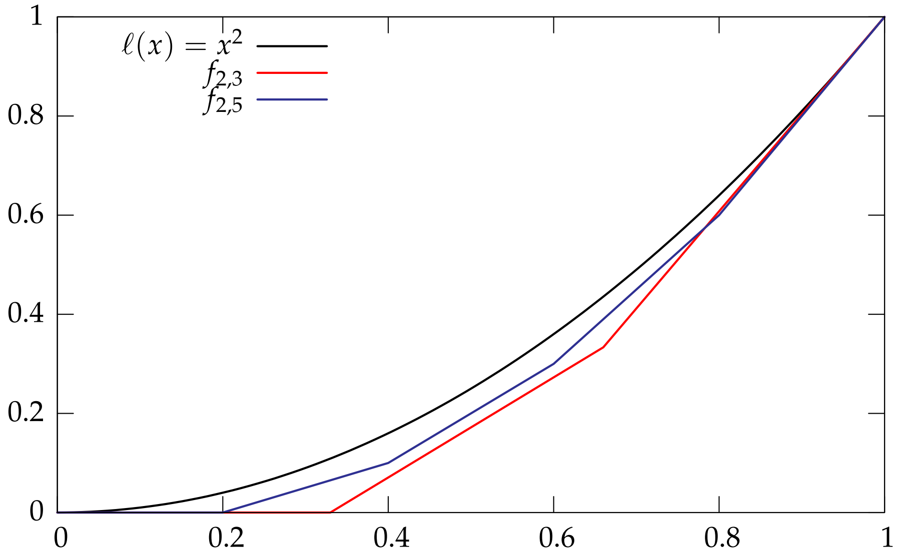

Figure A.1: The function $\ell(x) = x^2$ and the piecewise approximation $f_{2,n}$.

for any $x \in [0,1]$, $f_{n,d}(x) < \ell(x)$ for any function ℓ with elasticity d and $\ell(1) = 1$. To see this, note that between x_i and x_{i+1}, $f_{d,n}$ has slope

$$d \cdot \frac{f(x_{i+1})}{x_{i+1}} \quad ,$$

so the slope of $f_{d,n}$ is as least as large as the slope of any function with elasticity d, and property 1 follows by induction on i.

It remains to prove that $f_{d,n}$ satisfies property 2. For $i \geq d$ we can also write

$$f_{d,n}(x_i) = \prod_{j=i+1}^{n} \left(1 - \frac{d}{j}\right) \quad .$$

We see that this value is strictly positive since for $i \geq d$ every factor is strictly positive. We prove that, for any $i = d, \ldots, n$,

$$\int_0^{x_i} f_{d,n}(u)\, du \geq \left(\frac{1}{d+1} - \frac{d}{i}\right) \cdot x_i \cdot f_{d,n}(x_i) = \frac{x_i \cdot f_{d,n}(x_i)}{d+1} \cdot \left(1 - \frac{d(d+1)}{i}\right) .$$

Since this statement holds also for $i = n$, we obtain

$$\int_0^1 f_{d,n}(u)\, du \geq \frac{f_{d,n}(1)}{d+1} (1 - o(1)) \quad ,$$

as desired. The proof is by induction on i. Clearly, the statement holds for

$i = d$. Now assume that the statement holds for $i = j$. Then,

$$\begin{aligned}
\int_0^{x_{j+1}} f_{d,n}(u)\,du &\geq \int_0^{x_j} f_{d,n}(u)\,du + \frac{1}{n} \cdot f_{d,n}(x_j) \\
&\geq \left(\frac{1}{d+1} - \frac{d}{j}\right) \cdot x_j \cdot f_{d,n}(x_j) + \frac{1}{n} \cdot f_{d,n}(x_j) \\
&= \frac{x_{j+1} \cdot f_{d,n}(x_{j+1})}{d+1} \cdot \left(1 - \frac{(j+1) \cdot d(d+1) - d^3}{(j+1)^2}\right) \\
&\geq \frac{x_{j+1} \cdot f_{d,n}(x_{j+1})}{d+1} \cdot \left(1 - \frac{d(d+1)}{j+1}\right) \quad .
\end{aligned}$$

The first inequality is the definition of $f_{d,n}$, and the second is the induction hypothesis. The equality uses the definition of $x_j = j/n$ and $f_{d,n}(x_{j+1})$. This concludes the inductive argument. $\square$

Index

Bibliography

[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. In *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 613–622, 2006.

[2] Erwin Amann. *Evolutionäre Spieltheorie*. Physica-Verlag, 1999.

[3] Tom Anderson, Ratul Mahajan, Neil Spring, and David Wetherall. Rocketfuel: An ISP topology mapping engine. `http://www.cs.washington.edu/research/networking/rocketfuel/`.

[4] David Applegate and Edith Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 313–324, 2003.

[5] Sanjeev Arora, Yuval Rabani, and Umesh Vazirani. Simulating quadratic dynamical systems is PSPACE-complete. In *Proc. 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 459–467, 1994.

[6] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proc. 37th Annual ACM Symposium on Theory of Computing (STOC)*, 2005.

[7] Baruch Awerbuch, Rohit Khandekar, and Satish Rao. Distributed algorithms for multicommodity flow problems via approximate steepest descent framework. In *Proc. 18th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, 2007.

[8] Baruch Awerbuch and Robert D. Kleinberg. Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proc. 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 45–53, 2004.

[9] Baruch Awerbuch and Robert D. Kleinberg. Competitive collaborative learning. In *Proc. 18th Annual Conference on Learning Theory (COLT)*, pages 233–248, 2005.

[10] Paul Barford and Mark Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proc. ACM SIGMETRICS Conference*, pages 151–160, 1998.

[11] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[12] Petra Berenbrink, Tom Friedetzky, Leslie Ann Goldberg, Paul Goldberg, Zengjian Hu, and Russel Martin. Distributed selfish load balancing. In *Proc. 17th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, 2006.

[13] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computing: Numerical Methods*. Athena Scientific, 1989.

[14] Nam P. Bhatia and Giorgio P. Szergö. *Stability Theory for Dynamical Systems*. Springer-Verlag, 1970.

[15] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 45–52, Denver, Colorado, USA, July 2006. ACM.

[16] William Boyce. *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, 1965.

[17] Jianhua Cao and Christian Nyberg. An approximate analysis of load balancing using stale state information for servers in parallel. In *Proc. 2nd IASTED International Conference on Communications Internet and Information Technology*, November 2003.

[18] Zhiruo Cao, Zheng Wang, and Ellen W. Zegura. Performance of hashing-based schemes for Internet load balancing. In *Proc. Conference of the IEEE Communications Society (INFOCOM)*, pages 332–341, 2000.

[19] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proc. 37th Annual ACM Symposium on Theory of Computing (STOC)*, 2005.

[20] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? In *Proc. 4th ACM Conference on Electronic Commerce*, pages 98–107, 2003.

[21] Roberto Cominetti, José R. Correa, and Nicolás E. Stier Moses. Network games with atomic players. In *Proc. 33rd International EATCS Colloquium on Automata, Languages and Programming (ICALP)*, pages 525–536, 2006.

[22] Michael Dahlin. Interpreting stale load information. *IEEE Transactions on Parallel and Distributed Systems*, 11(10):1033–1047, 2001.

[23] E. Dekel and S. Scotchmer. On the evolution of optimizing behavior. *Journal of Economic Theory*, 57:392–406, 1992.

[24] Rodney D. Driver. Existence and stability of solutions of a delay-differential system. *Arch. Rational Mech. Anal.*, 10:401–426, 1962.

[25] Anwar Elwalid, Cheng Jin, Steven H. Low, and Indra Widjaja. MATE: MPLS adaptive traffic engineering. In *Proc. Conference of the IEEE Communications Society (INFOCOM)*, pages 1300–1309, 2001.

[26] Kousha Etessami and Andreas Lochbihler. The computational complexity of evolutionarily stable strategies. Technical Report TR04-055, Ellectronic Colloquium on Computational Complexity, 2004.

[27] Eyal Even-Dar, Michael Kearns, Yishay Mansour, and Jennifer Wortman. Regret to the best vs. regret to the average. In *Proc. 20th Annual Conference on Learning Theory (COLT)*, 2007.

[28] Eyal Even-Dar and Yishay Mansour. Fast convergence of selfish rerouting. In *Proc. 16th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 772–781, 2005.

[29] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

[30] Anja Feldmann, Anna Gilbert, Polly Huang, and Walter Willinger. Dynamics of IP traffic: A study of the role of variability and the impact of control. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 163–174, 1999.

[31] Simon Fischer, Nils Kammenhuber, and Anja Feldmann. REPLEX — Dynamic traffic engineering based on Wardrop routing policies. In

Proc. 2nd Conference on Future Networking Technologies (CoNext), pages 6–17, Lisboa, Portugal, December 2006.

[32] Simon Fischer, Harald Räcke, and Berthold Vöcking. Fast convergence to Wardrop equilibria by adaptive sampling methods. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 653–662, Seattle, WA, USA, May 2006. ACM.

[33] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. In Susanne Albers and Tomasz Radzik, editors, *Proc. 12th Annual European Symp. on Algorithms (ESA)*, number 3221 in Lecture Notes in Computer Science, pages 323–334, Bergen, Norway, September 2004. Springer-Verlag.

[34] Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. In Marcos Kawazoe Aguilera and James Aspnes, editors, *Proc. 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 276–283, Las Vegas, NV, USA, July 2005. ACM.

[35] Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. In *Proc. 31st International EATCS Colloquium on Automata, Languages and Programming (ICALP)*, pages 544–554, 2004.

[36] Bernard Fortz and Mikkel Thorup. Internet traffic engineering by optimizing OSPF weights. In *Proc. Conference of the IEEE Communications Society (INFOCOM)*, 2000.

[37] Bernard Fortz and Mikkel Thorup. Optimizing OSPF/IS-IS weights in a changing world. *IEEE Journal on Selected Areas in Communications*, 20(4), 2002.

[38] Vasyl V. Gafiychuk and Anatolyi K. Prykarpatsky. Replicator-mutator evolutionary dynamics. *Journal of Nonlinear Mathematical Physics*, 11(3):350–360, 2004.

[39] Wang Hao and Mabo R. Ito. Dynamics of load-sensitive adaptive routing. In *Proc. IEEE International Conference on Communications (ICC)*, 2005.

[40] Alain B. Haurie and Patrice Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.

[41] Ara Hayrapetyan, Éva Tardos, and Tom Wexler. The effect of collusion in congestion games. In *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 89–98, 2006.

[42] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[43] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, 2003.

[44] Srikanth Kandula, Dina Katabi, Bruce Davie, and Anna Charny. Walking the tightrope: Responsive yet stable traffic engineering. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 253–264, 2005.

[45] D. Katz, K. Kompella, and D. Yeung. Traffic engineering extensions to OSPF version 2. Internet Draft, `http://ietfreport.isoc.org/all-ids/draft-katz-yeung-ospf-traffic-10.txt-28939.txt`, 2003.

[46] Leonid G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Math. Dokl*, 20:191–194, 1997.

[47] Atul Khanna and John A. Zinky. The revised ARPANET routing metric. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 45–56. ACM, September 1998.

[48] Balachander Krishnamurthy and Jennifer Rexford. *Web Protocols and Practice*. Addison-Wesley, 2001.

[49] James F. Kurose and Keith W. Ross. *Computer Networking, a top down approach featuring the Internet, 3rd ed.* Addison-Wesley Longman, 2004.

[50] Michael Laor and Lior Gendel. The effect of packet reordering in a backbone link on application throughput. *IEEE Network*, 2002.

[51] Yong Liu and A. L. Narasimha Reddy. Route optimization among a group of multihomed stub networks. In *Proc. IEEE GLOBECOM*, 2006.

[52] Michael Mitzenmacher. How useful is old information? *IEEE Transactions on Parallel and Distributed Systems*, 11(1):6–20, January 2000.

[53] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.

[54] Randolph Nelson. *Probability, stochastic processes, and queueing theory. The mathematics of computer performance modeling.* Springer-Verlag, New York, 1995.

[55] Noam Nisan. A note on the computational complexity of evolutionary stable strategies. Technical Report TR06-076, Electronic Colloquium on Computaional Complexity, 2006.

[56] Martin J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2003.

[57] Kihong Park and Walter Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. Wiley-Interscience, 2000.

[58] Yuri Rabinovich, Alistair Sinclair, and Avi Wigderson. Quadratic dynamical systems. In *Proc. 33rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 304–313, 1992.

[59] Vivek Raghunathan and P. R. Kumar. A Wardrop routing protocol for ad hoc wireless networks. In *Proc. 43rd IEEE Conference on Decision and Control (CDC)*, pages 14–17, 2004.

[60] Vivek Raghunathan and P. R. Kumar. Issues in Wardrop routing in wireless networks. In *Proc. 1st IEEE Wireless Internet Conference (WICON)*, pages 34–41, 2005.

[61] Renesys. The SSFNet network simulator. `http://www.ssfnet.org/`.

[62] Jennifer Rexford. *Handbook of Optimization in Telecommunications*, chapter Route optimization in IP networks. Kluwer Academic Publishers, 2005.

[63] Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[64] Matthew Roughan, Mikkel Thorup, and Yin Zhang. Traffic engineering with estimated traffic matrices. In *Proc. 3rd ACM SIGCOMM Conference on Internet Measurement (IMC)*, pages 248–258, 2003.

[65] Tim Roughgarden. Stackelberg scheduling strategies. In *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 104–113, 2001.

[66] Tim Roughgarden. How unfair is optimal routing? In *Proc. 13th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 203–204, 2002.

[67] Tim Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–367, September 2003.

[68] Tim Roughgarden. The maximum latency of selfish routing. In *Proc. 15th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 973–974, 2005.

[69] Tim Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, May 2005.

[70] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[71] Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.

[72] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391, 1992.

[73] Stefan Savage, Andy Collins, Eric Hoffman, John Snell, and Thomas E. Anderson. The end-to-end effects of internet path selection. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 289–299, Boston, MA, 1999.

[74] Alexander Schrijver. *Combinatorial Optimisation*, volume A. Springer-Verlag, 2003.

[75] Nataliya Skrypnyuk. Load-sensitive routing. Master's thesis, Technische Universität München, 2006.

[76] Neil Spring, Ratul Mahajan, and David Wetherall. Measuring ISP topologies with Rocketfuel. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 133–145, Pittsburgh, PA, USA, August 2002. ACM.

[77] Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1–2):145–156, 1978.

[78] Renata Teixeira, Nick Duffield, Jennifer Rexford, and Matthew Roughan. Traffic matrix reloaded: Impact of routing changes. In *Proc. 6th Passive and Active Network Measurement (PAM)*, pages 251–264, 2005.

[79] TOTEM—TOolbox for Traffic Engineering Methods. `http://totem.run.montefiore.ulg.ac.be/`.

[80] Fernando Vega-Redondo. *Economics and the Theory of Games*. Cambridge University Press, 2003.

[81] Christian Vollmert. A Web workload generator for the SSFNet network simulator. Bachelor's thesis, Technische Universität München, 2004.

[82] Jörg Wallerich. Design and implementation of a WWW workload generator for the NS-2 network simulator. Master's thesis, Universität des Saarlandes, 2001.

[83] John Glen Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.

[84] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.

[85] Wen Xu and Jennifer Rexford. MIRO: Multi-path interdomain routing. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, 2006.

[86] Chun Zhang, James F. Kurose, Donald F. Towsley, Zihui Ge, and Yong Liu. Optimal routing with multiple traffic matrices: Tradeoff between average case and worst case performance. In *Proc. 13th International Conference on Network Protocols (ICNP)*, 2005.

[87] Yin Zhang, Matthew Roughan, Carsten Lund, and David Donoho. An information-theoretic approach to traffic matrix estimation. In *Proc. ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, 2003.

Index

Curriculum Vitae

Simon Fischer
Rosfeld 76
52074 Aachen

geboren am 26. Juli 1978 in Hagen
verheiratet, eine Tochter

Ausbildung

1998 – 2003	Studium der Informatik mit Nebenfach Physik Universität Dortmund Abschluss: Diplom
Sommer 2001	Sommerstudent am Europäischen Kernforschungszentrum (CERN) in Genf
1988 – 1997	Geschwister-Scholl-Gymnasium, Wetter (Ruhr) Abschluss: Abitur
1984 – 1988	Gemeinschaftsgrundschule Esborn, Wetter (Ruhr)

Beruflicher Werdegang

03/2005 – jetzt	Wissenschaftlicher Mitarbeiter am Lehrstuhl für Algorithmen und Komplexität der RWTH Aachen
01/2004 – 02/2005	Wissenschaftlicher Mitarbeiter am Lehrstuhl für Algorithmen und Komplexität der Universität Dortmund
10/2000 - 12/2003	Studentische Hilfskraft am Lehrstuhl für künstliche Intelligenz der Universität Dortmund

Zivildienst

07/1997 – 07/1998	Individuelle Schwerstbehindertenbetreuung

Simon Fischer