Part III

Equilibrium and Learning in Traffic Networks

Roberto Cominetti

Foreword

To Cecilia, Andrés and Pablo

These notes provide a brief introduction to some of the mathematical aspects involved in the study of equilibrium models for congested traffic networks. They were written to support a series of lectures delivered by the author at the Centre de Recerca Matemàtica (Barcelona, Spain) in July 2009.

Time and space limitations imposed a choice of topics to be covered and the level of detail in their treatment. After reviewing the classical concepts of Wardrop and stochastic user equilibrium, we provide a detailed treatment of the notion of Markovian traffic equilibrium. From these static equilibrium concepts, we move on to describe recent work on an adaptive procedure that models travel behavior, and its asymptotic convergence towards equilibrium. While this choice reflects the author's bias, I hope it serves as an introduction to the subject and to motivate a deeper study of an area that has still many interesting open questions.

I would like to express my gratitude to the organizers of the Advanced Courses at CRM, and particularly to Aris Daniilidis and Juan Enrique Martínez-Legaz, for their invitation and for making this course possible. I am also deeply indebted to Vianney Perchet and Guillaume Vigeral for their careful and thoughtful revision that greatly contributed to improve these notes.

Santiago de Chile, September 2010

Roberto Cominetti

Introduction and Overview

Congestion is a common characteristic of large urban areas, especially at peak hours when transport demands approach the saturation capacities of the roads. Urban planning and traffic management policies require quantitative models to forecast traffic in order to evaluate and compare alternative designs. These notes provide a brief introduction to some of the mathematical tools involved in the modeling of traffic flows in congested networks.

Traffic is often described as a steady state that emerges from the adaptive behavior of selfish drivers who strive to minimize travel times while competing for limited road capacity. Such equilibrium models provide a static description of how traffic demands flow through a congested network. In 1952, Wardrop [76] proposed a first model using continuous variables to represent aggregate flows. Drivers are assumed to behave rationally by selecting shortest paths according to the prevailing traffic conditions, while congestion is modeled by travel times that increase with the total flow carried by each route. In this setting, any given traffic pattern induces route loads and travel times that determine in turn which routes are optimal. An equilibrium is defined as a consistent traffic pattern for which the routes that are actually used are the shortest ones according to the induced travel times. Soon after the model was introduced, Beckman et al. [11] realized that these equilibria could be characterized as the optimal solutions of an equivalent convex minimization problem. In the 1980s, these results were supplemented by Daganzo [28] and Fukushima [43] who described a dual characterization that led to alternative numerical methods to compute the equilibrium. In Chapter 1 we present a brief overview of *Wardrop equilibrium* and its characterization.

The relatively poor agreement between the traffic patterns predicted by Wardrop equilibrium and the ones observed in real networks led to a critical revision of the assumptions that supported the model. One criticism focused on the assumption of homogeneity in driver behavior. As a matter of fact, drivers' decisions are subject to random fluctuations that arise from the difficulty in discriminating routes with similar costs, as well as from the intrinsic variability of travel times that induce different perceptions of the routes depending on the individual past experience. The idea of a random selection of routes is captured by discrete choice models based on random utility theory, leading to the notion of *stochastic user equilibrium* (SUE) [27, 28, 30, 35, 58, 62, 70, 74]. In Chapter 1 we give a short overview of this concept. More detailed accounts can be found in the book by Ben-Akiva & Lerman [16] and the surveys by Florian & Hearn [36, 37].

As the network size increases, the number of routes grows exponentially and the assumption that drivers are able to compare all the routes becomes less and less plausible. The exponential growth also makes route-based models computationally intractable, so that methods to circumvent path enumeration were developed in [2, 3, 17, 51, 55]. These ideas evolved into the notion of *Markovian equilibrium* [9] which looks at route choice as a stochastic dynamic process: drivers proceed towards their destination by a sequential process of arc selection using a discrete choice model at every intermediate node in their trip. The route is no longer fixed at the origin but it is the outcome of this sequential process, so that driver movements are governed by a Markov chain and the network flows are the corresponding invariant measures. In Chapter 2 we discuss this model in detail.

Equilibrium models assume implicitly the existence of an underlying mechanism in travel behavior that stabilizes traffic flows at a steady state. However, the models are stated directly as equilibrium equations at an aggregate population level and are not tied to a specific adaptive mechanism of individual drivers. Empirical evidence of adaptive behavior based on experiments and simulations has been reported in the literature [8, 31, 49, 57, 68, 77], though it was observed that the resulting steady states may differ from the standard notions of equilibria. Also, several continuous time dynamics describing plausible adaptive mechanisms that converge to Wardrop equilibrium were studied in [41, 67, 73] while a class of finite-lag discrete time adjustment procedures was considered in [23, 24, 29, 47]. However, these dynamics are again of an aggregate nature and are not explicitly tied to the behavior of individual drivers. An alternative to Wardrop's model in which drivers are considered as individual players in a game was studied by Rosenthal [64]. The main result established the existence of a Nash equilibrium in pure strategies by exploiting a potential function which is a discrete analog of Beckman et al. [11]. This motivated a number of extensions to the so-called *congestion games* and the more general class of *potential games*. Although we will not review these contributions, in Chapter 3 we reconsider the atomic framework with finitely many drivers but in a dynamic setting that models their adaptive behavior.

Learning and adaptation are fundamental issues in the study of repeated games with boundedly rational players. They have been intensively explored in the last decades [42, 79], though most results apply to games with a small number of players. The most prominent adaptive procedure is *fictitious play*, which assumes that at each stage players choose a best reply to the observed empirical distribution of past moves by their opponents [22, 63]. The assumption that players are able to record the past moves of their opponents is very stringent for games involving many players with limited observation capacity. A milder assumption is that players observe only the outcome vector, namely the payoff obtained at every stage and the payoff that would have resulted if a different move had been played. Procedures such as *no-regret* [44, 45], *exponential weight* [40], and *calibration* [38], deal with such limited information contexts assuming that players adjust their behavior based on payoff statistics. Eventually, adaptation leads to configurations where no player regrets the choices he makes. Although these procedures are flexible and robust, the underlying rationality may still be too demanding for games with a large number of strategies and poorly informed players. This is the case for traffic where a multitude of small players make routing decisions with little information about the strategies of other drivers nor the actual congestion in the network.

A simpler adaptive rule that relies only on the sequence of realized moves and payoffs is *reinforcement* [7, 19, 33], where players select moves proportionally to the cumulative payoff of each alternative. In Chapter 3 we describe a similar approach to model the adaptive behavior of drivers in a simple network with parallel links: each player observes only the travel time of the specific route chosen on any given day, and future decisions are adjusted based on past observations. Specifically, each player has a prior estimate of the average payoff of each route and makes a decision based on this rough information using a random choice rule. The payoff of the chosen route is then observed and is used to update the perception for that particular move. This procedure is repeated day after day, generating a discrete time stochastic process called the *learning process*. Since travel times depend on the congestion of routes imposed collectively by all the players' decisions, the process progressively reveals to each player the congestion conditions on all the routes. In the long run these dynamics lead the system to coordinate on a steady state that can be characterized as a Nash equilibrium for a particular limit game. Convergence depends on a viscosity parameter that represents the amount of noise in players' choices. If noise is large enough, the adaptive dynamics have a unique global attractor which almost surely attracts the sequences generated by the learning process. This approach proceeds bottom-up: a simple and explicit discrete time stochastic model for individual behavior gives rise to an associated continuous time deterministic dynamics which leads ultimately to an equilibrium of a particular limit game. The equilibrium is not postulated a priori but it is derived from basic assumptions on player behavior, providing a microeconomic foundation for equilibrium.