Chapter 1

Introduction

A standard problem in network optimization is to find a routing of traffic demands from sources to destinations using a given network infrastructure. This problem is also known as a *multicommodity flow problem* or *traffic assignment problem*. A cost is associated with every arc of the network, which is usually a function of the arc flow. Most of the optimization models assume a central planer controlling the whole system and determining the best possible routing accordingly. This routing is also called the *system optimum*. For many real-world applications, this assumption is problematic with respect to several aspects: It is not always assured that (i) a central planer has access to the necessary information (*information problem*), (ii) the best possible solution is efficiently computable even if the needed information is accessible (*complexity problem*), (iii) the individual traffic sources agree to a proposed solution (*implementation problem*).

A tremendous amount of effort has been invested in designing efficient routing algorithms to cope with the above problems. As an example for the information problem, consider a traffic assignment problem, where demands arrive sequentially in time in an online fashion. An algorithm that routes these demands without knowledge about future demands is called an *online algorithm*. The common theoretical concept to evaluate the efficiency of an online algorithm is based on *competitive analysis* coming from the online optimization field. An online algorithm is called competitive, if its cost is never larger than a constant factor times the cost of an optimal offline solution. Another research area that covers the complexity problem is concerned with deriving efficient algorithms for solving (NP-hard) optimization problems. Of particular interest is the notion of an *approximation ratio* for a heuristic to solve such optimization problems. The approximation ratio is defined as the largest ratio of the objective value obtained by the solution of the heuristic and that of an optimal solution. The implementation problem can be analyzed within the algorithmic game theory field. Here, one tries to quantify the efficiency loss caused by selfish users compared to the system optimum. The cost of this lack of coordination has been coined "price of anarchy" by Koutsoupias and Papadimitriou in [62]. While the approximation ratio and competitive ratio measure the worst case loss in solution quality due to insufficient computing power and information, respectively, the price of anarchy measures the worst case loss due to insufficient ability to control and coordinate the actions of selfish individuals.

All three issues are exemplified by several practical applications that have motivated the topics covered in this thesis. For instance, billions of packets traverse the world wide web along routes that are decided on by Internet routing protocols. This routing is done in an online fashion without knowledge about future traffic changes. The size of the Internet and the heterogeneity of Internet applications contribute to the computational complexity of finding the best possible routing. Furthermore, a centrally coordinated implementation usually contradicts security requirements of Internet users. Another example for the implementation problem is the road traffic network, where the majority of traffic follows routes that are chosen based on selfish interests of the individuals. It is well known that some users would have to take long detours in a system optimal routing, which makes such a solution unattractive for the affected users.

The main topic of this thesis is to study multicommodity flow problems that exhibit a combination of the afore mentioned three problems. In particular, we focus on online multicommodity routing problems, selfish routing problems, and a combination of these two problems. Thus, the theoretical concepts that we use to analyze the corresponding routing patterns stems from competitive analysis and bounding the price of anarchy. Indirectly, these concepts also provide an approximation ratio, since a solution produced by an online algorithm or a solution produced by selfish individuals constitutes an approximation for the optimal solution of an optimization problem.

1.1 Online Multicommodity Flow Problems

In the first part of this thesis, we study online multicommodity routing problems, where demands have to be routed sequentially in a network. The cost of a flow is determined by dynamic load dependent price functions on links. We make four crucial assumptions: (i) demands for commodities are revealed in an *online* fashion and have to be routed immediately; (ii) demands can be split along several paths; (iii) once a demand is routed, no rerouting is allowed; (iv) the routing cost on an arc is given by the integral over the arc flow with respect to the corresponding price function. Since at the time of routing a commodity, future demands are not known, this yields an online optimization problem that we call the *Online Multicommodity Routing Problem*.

This problem arises in an inter-domain resource market in which multiple service providers offer network resources (capacity) to enable Internet traffic with specific Quality of Service (QoS) constraints, see for example Yahaya and Suda [88] and Yahaya, Harks, and Suda [90]. In such a market, each service provider advertises prices for resources that he wants to sell. We assume that prices are determined by load dependent price functions. Buying providers reserve capacity along paths to route demand (coming from own customers) from sources to destinations via domains of other providers. The routing of a demand along paths is fixed by establishing a contract between the source domain and all domains along the paths. Prices in the market, however, are only valid for a predefined bundle size, that is, after the routing of flow with this bundle size, the arc prices are updated. In the limiting case, where the bundle size tends to zero, the routing cost on an arc is given by the integral over the arc flow with respect to the corresponding price function.

Contributions (Chapter 3)

We investigate such multicommodity online routing problems and allow for arbitrary continuous and nondecreasing load dependent price functions defining the routing costs. We investigate a greedy online algorithm, called SEQ, for this setting and investigate, in which cases this algorithm is competitive. Our main finding is that for polynomial price functions with nonnegative coefficients, the competitive ratio of SEQ can be bounded by a constant factor that only depends on the maximum degree of the polynomials but is independent of the network topology and demand sequence. For the single-source single-destination case, we show that this algorithm is optimal. Without restrictions on the price functions and network, no algorithm is competitive. We also investigate a variant in which the demands have to be routed unsplittably. In this case, the offline problem is NP-hard. As in the splittable case, in general there exists no competitive deterministic online algorithm. For linear price functions, any deterministic online algorithm has a competitive ratio of at least 2. Finally, we present a computational study for unsplittable routings for a realistic network topology and stochastically generated demands. Our empirical findings state that the efficiency loss is significantly smaller in this case compared to the provable upper bounds for the splittable online routing SEQ. The online algorithm SEQ and the ONLINEMCRP can be viewed as a first step towards a methodology for analyzing the efficiency of general inter-domain routing strategies. These results are presented in Chapter 3.

1.2 Network Games

Second, we study the impact of selfish behavior on social welfare in network games. We are interested in the degradation of system performance if players select routes based on selfish interest. Consider a network of arcs that are used by individuals to route demand from sources to destinations. A common approach is to model congestion on arcs by nondecreasing latency functions mapping the flow on an arc to the time needed to traverse this arc. Since individuals share the same network, congestion effects on arcs generate interdependencies between the routing decisions. In this regard, non-cooperative game theory provides the appropriate concepts to analyze such interdependencies. In a non-cooperative game, players compete for shared resources and the utility of each individual player depends on the number of players that choose the same or some overlapping strategy, see Rosen [78]. In the network routing context, the strategies correspond to the available routes and the utility of a player is its total travel time. A classical approach to describe the outcome of a non-cooperative game is to analyze an equilibrium situation. The most popular notion of such an equilibrium is the *Nash equilibrium*: a stable point from which no individual has an incentive to deviate unilaterally. In nonatomic network games in which a single individual player has only a negligible impact on the travel time of others, Wardrop [87] characterized such an equilibrium in his first principle as follows. All path flows between a single source and a single destination have equal latency. A Wardrop equilibrium can be interpreted as a Nash equilibrium in this case.

A fundamental question that has already been raised in 1920 by Pigou [75] and later on in the 1950's by Wardrop [87] and Beckmann, McGuire and Winsten [11] is the following: How efficient is the performance of a Nash equilibrium compared to the best possible outcome? As already noted, the cost of this lack of coordination is called *price of anarchy*. For the Wardrop traffic model, Roughgarden and Tardos [84] proved in a seminal paper that the total travel time of a flow at Nash equilibrium does not deviate too much from the minimum total travel time. In particular, they proved that the price of anarchy is bounded by 4/3 provided affine linear latency functions are considered. By introducing the so called *anarchy value* $\alpha(\mathcal{L})$ for a class $\mathcal{L}$ of latency functions, Roughgarden [81] proved the first tight bounds on the price of anarchy for general polynomial latency functions. Correa, Schulz, and Stier-Moses [24] introduced a different parameter $\beta(\mathcal{L})$ that allows to relax some previous assumptions on allowable latency functions. They proved that their bound implies all bounds of Roughgarden by using the relation $\alpha(\mathcal{L}) = \big(1 - \beta(\mathcal{L})\big)^{-1}$.

Even though we have just argued that the outcome of a Nash equilibrium is not too inefficient, there has been a recent trend towards using route guidance devices for improving the individual travel time. The current position of each driver is determined via the *Global Positioning System* (GPS) at the beginning of a trip. A central computer calculates then an "optimal" route for this trip based on digital maps and on available knowledge of traffic congestion on the streets. In game theoretic language, a route guidance operator is an *atomic player* since a significant (non negligible) part of the entire demand is controlled. Roughgarden [83] and Correa, Schulz, and Stier-Moses [25] claimed that the price of anarchy in an atomic network game does not exceed that of the corresponding nonatomic one. Interestingly, this turned out to be wrong, as reported by Cominetti, Correa, and Stier-Moses in [23]. Based on the work of Catoni and Pallotino [19], they presented an example in which the price of anarchy in a network game with atomic players is larger than that of the corresponding nonatomic game. Moreover, they showed that the cost for an atomic player may even increase compared to the nonatomic game. Such a counter-intuitive phenomenon can also arise from the perspective of single individuals: a nonatomic player competing with an atomic player may face lower cost compared to the situation in which the atomic player is replaced by nonatomic ones. Cominetti, Correa, and Stier-Moses showed that the price of anarchy for the atomic network game can be bounded for special latency functions. In particular, they proved upper bounds of $1.5, 2.56$, and 7.83 on the price of anarchy for affine linear, squared, and cubic latency functions with nonnegative coefficients, respectively. For polynomials with nonnegative coefficients and higher degree, their approach fails to generate upper bounds on the price of anarchy.

Contributions (Chapter 4)

For network games with nonatomic players, we introduce the value $\omega(\mathcal{L}, \lambda)$ for bounding the price of anarchy. This value generalizes the anarchy value $\alpha(\mathcal{L})$ and the value $\beta(\mathcal{L})$. Using our value, we reprove the existing tight bounds on the price of anarchy and present a novel proof for monomial latency functions showing that the price of anarchy is one in this case.

For network games with atomic players, we improve all previously known bounds for polynomial latency functions with nonnegative coefficients, except for affine linear latency functions. These results are presented in Chapter 4.

1.3 Online Network Games

Combining the online aspect with selfish behavior of individuals, we investigate an online routing problem called *online network games*. In this problem, we assume a sequence of network games $\sigma = (1, \ldots, n)$ that are released consecutively in time in an online fashion. A network game is characterized by a set of commodities that have to be routed in a given network. Arcs in the network are equipped with load dependent latency functions defining the routing cost. By the time of routing commodities of game i, future games $i + 1, \ldots, n$ are not known. We further assume that once commodities of a game are routed, this routing remains fixed, that is, the routings are irrevocable. We analyze two online algorithms, called NSeqNash and ASeqNash. These algorithms produce a flow consisting of a sequence of Nash equilibria for the corresponding games with nonatomic and atomic players, respectively. As usual, we analyze the efficiency of an online algorithm in terms of competitive analysis.

The online variant of network games is motivated by the application of selfish routing to the source routing concept in telecommunication networks, see Qiu, Yang, Zhang, and Shenker [76] and Friedman [42] for an engineering perspective and Roughgarden [80] and Altman, Basar, Jimenez, and Shimkin [5] for a theoretical perspective on this topic. In the source routing model, sources are responsible for selecting paths to route data to the corresponding destination. The arcs in the network advertise their current status that is based on the current congestion situation. If the costs on arcs correspond to the expected delay, minimum cost routing is a natural goal for real-time applications.

As described in the last section, the main focus of the line of research that studies source routing is to quantify the price of anarchy. Here, one assumption is crucial: if the traffic matrix changes, all sources may possibly change their routes and form a new equilibrium. This assumption, however, has some important implications: Each source would have to *continuously* maintain the current state of all available routes, which in turn introduces additional traffic overhead by continuously signaling this needed information. Furthermore, frequent rerouting attempts during data transmission may not only produce transient load oscillations but may also interfere with the widely used congestion control protocol TCP that controls the data rate, as reported by La, Walrand, and Anantharam in [63]. For these reasons, frequent rerouting attempts in reaction to traffic changes in the network are not necessarily beneficial. Time critical applications, such as Internet telephony or video streaming may suffer severe performance degradation.

Contributions (Chapter 5)

In this regard, we propose and investigate a new model, called *Online Network Games*, in which sources starting at the same time select their routes *only* during connection setup phase. We then study the extreme case in which flows fix their routing decisions once they are at equilibrium. Thus, continuously gathering information about the state of available routes is not necessary after this initial routing game. Relying on competitive analysis, we analyze online algorithms that produce a flow that is at Nash equilibrium for every game out of a sequence of games. The cost function is given by the total travel cost after all games have been played. Our main result states that for polynomial latency functions with nonnegative coefficients, the competitive ratio of both NSeqNash, andASeqNash can be bounded by a constant factor, which depends on the maximum degree. This result holds independently of the network topology or game sequence. We also prove lower bounds. In particular, we show that for a sequence of two network games and affine linear latency functions, our upper bound for the NSeqNash is tight. Furthermore, we prove for a given sequence of games and parallel arcs that the competitive ratio of the online algorithm NSeqNash does not exceed the price of anarchy of a complementary nonatomic network game in which all commodities of the sequence of games are considered at the same time.

1.4 Thesis Organization

After describing the motivation and background for this thesis in Chapter 1, we present in Chapter 2 the basic concept of competitive analysis in online optimization. In Chapter 3, we present the framework OnlineMCRP in which we study the online algorithm Seq. In Chapter 4 we focus on network games with nonatomic and atomic players, respectively. Finally, we combine network games with online aspects in Chapter 5.

We note that the "Contribution and Chapter Outline" section at the beginning of Chapter 3, 4, and 5 gives an overview and road map about the results presented in that chapter. We further recommend that Chapter 4 is read prior to Chapter 5. Except for introducing the notation for multicommodity flow problems, Chapter 4 and 5 can be read independently from Chapter 3.