CCIT Report #461
January 2004

The Price of Anarchy in Networks with Bottleneck Objectives

Ron Banner and Ariel Orda
Department of Electrical Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
{banner@tx, ariel@ee}.technion.ac.il

Total number of pages (including Appendix): 13

Abstract—**We study the price of anarchy when performance is dictated by the worst (bottleneck) element. We are given a network, finitely many users, each associated with a positive flow demand, and a load-dependent performance function for each network element; the network objective is to route traffic such that the performance of the worst (i.e., bottleneck) element in the network is optimized. In the absence of regulation by some central authority, we assume that each user routes its traffic selfishly i.e., through paths that optimize the *user's* performance, in terms of bottleneck elements. We prove the existence of a Nash equilibrium, considering two routing scenarios, namely when a user can split its traffic over more than one path and when it cannot. Then, we quantify the price of anarchy in both scenarios. Specifically, we show that, when users are allowed to split their traffic, *anarchy comes at no price*. On the other hand, we show that, if each user is limited to a single path, the price of anarchy is *unbounded*. Finally, we turn to consider the case where the network objective is *additive* i.e., minimizing the sum of all link performance functions, while users still optimize the performance of their bottleneck elements. For this case, and for users that can split their traffic, we show that the price of anarchy is at most the number of network links. We then delineate a possible application of this result for the case where both the network and the users consider *additive* objectives.**

***Keywords*- bottleneck & additive metrics, Nash equilibrium, price of anarchy/coordination factor, selfish routing, unregulated traffic.**

I. INTRODUCTION

Traditional computer networks were designed and operated with systemwide optimization in mind. Accordingly, the actions of the network users were determined so as to optimize the overall network performance. Consequently, users would often find themselves sacrificing some of their own performance for the sake of the entire network. In recent years it has been recognized that systemwide optimization may be an impractical paradigm for the control of modern networking configurations [1],[12],[16],[27]. Indeed, control decisions in large scale networks are often made by each user independently, according to its own individual performance objectives. Such networks are henceforth called *noncooperative*, and Game Theory [17] provides the systematic framework to study and understand their behavior.

Game theoretic models have been employed in the context of flow control [1],[12],[27] routing [16],[20],[21] and bandwidth allocation [14]. These studies mainly investigated, the structure of network operating points i.e., the Nash equilibria of the respective games. Such equilibria are inherently inefficient [11] and, in general, exhibit suboptimal network performance. In order to understand this phenomenon, an investigation of the performance ratio between the worst possible Nash equilibrium and the social (i.e., overall) optimum was initiated in [13]. This ratio, termed the *price of anarchy* (also: *coordination factor*), was first investigated by [10],[13],[15] for routing problems in which a set of users send traffic along a set of parallel links with linear cost functions. A more general framework for general topologies was later considered in [19],[20],[21]; in those studies, the cost of each link was a load-dependent latency function, and each network user chose a minimum-latency path while controlling a negligible fraction of the overall traffic; the network objective was to minimize total latency.

The above studies solely focused on *additive metrics* i.e., the case where performance is determined by the *sum* of link cost functions. Another fundamental case is that of *bottleneck metrics*, in which network performance is determined by the *worst* component (link) in the network. Bottleneck objectives (also known as Max-Min or Min-Max objectives) are of major practical importance. For example, a commonly used objective for traffic engineering is to minimize the utilization of the most utilized link in the network, in order to move away traffic from congested hot spots to less utilized parts of the network [2],[26]. Another example is when the performance goal is to maximize the ability to accommodate momentary traffic bursts by maximizing the minimum residual capacity (or