

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JAN 2004	2. REPORT TYPE	3. DATES COVERED 00-00-2004 to 00-00-2004			
4. TITLE AND SUBTITLE The Maximum Latency of Selfish Routing		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cornell University, Department of Computer Science, Ithaca, NY, 14853		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES ACM-SIAM Symposium on Discrete Algorithms (SODA04), 11-13 Jan 2004, New Orleans, LA. U.S. Government or Federal Rights License					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

The Maximum Latency of Selfish Routing

Tim Roughgarden*

1 Introduction

We study a model of “selfish routing” first studied in a theoretical computer science context by Roughgarden and Tardos [3]. In this model, we are given a directed network in which each edge possesses a continuous, non-decreasing latency function that describes the common delay experienced by all traffic on the edge as a function of the edge congestion. We assume that each network user acts selfishly and routes itself from its source to its desired destination to minimize the latency experienced, given the network congestion due to the other users. As in most earlier works, we assume that the traffic comprises a large population of users, so that the actions of a single individual have negligible effect on network congestion.

There are many senses in which a “selfish” assignment of traffic to paths—a *Nash equilibrium*—is inefficient. For example, it is well known that selfish routing does not minimize the average or maximum latency experienced by traffic. Previous papers [2, 3] have precisely quantified how much the average latency of traffic at Nash equilibrium can exceed that of an optimal routing—the “price of anarchy” with respect to the average latency.

Less is known about the price of anarchy of selfish routing relative to the *maximum* latency incurred by network traffic, though some bounds on this quantity for single-commodity networks are implicit in previous work [1, 2, 3]. Corollaries of these papers include nearly tight constant bounds on the price of anarchy w.r.t. restricted classes of allowable latency functions (such as bounded-degree polynomials), and a lower bound of $\lfloor n/2 \rfloor$ for networks with n vertices and arbitrary continuous, nondecreasing latency functions. Weitz [4] was the first to explicitly consider the price of anarchy of selfish routing relative to the maximum latency, and he proved that the price of anarchy is at least approximately $n/2$ in multicommodity networks with linear latency functions.

In this note, we prove that for single-commodity networks with n vertices and arbitrary continuous, nondecreasing latency functions, the price of anarchy is precisely $n - 1$. We thus give both the first finite upper bound, and a new, optimal lower bound. The finite upper bound stands in contrast to the price of anarchy relative to the average latency, which is unbounded even for two-node two-link networks [3]. We also give two conjectures (but no results) on the price of anarchy for multicommodity networks.

2 Preliminaries

The Model. We consider a directed network $G = (V, E)$ with vertex set V , edge set E , and source-destination pairs $\{s_1, t_1\}, \dots, \{s_k, t_k\}$. We denote the set of s_i - t_i paths by \mathcal{P}_i , assume that each such set is nonempty, and define $\mathcal{P} = \cup_i \mathcal{P}_i$. A *flow* is a function $f : \mathcal{P} \rightarrow \mathcal{R}^+$; for a fixed flow f we define the load $f_e = \sum_{P: e \in P} f_P$. With respect to a finite and positive *traffic rate vector* r , a flow f is said to be *feasible* if $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for all $i \in \{1, 2, \dots, k\}$. Each edge $e \in E$ is given a load-dependent *latency* that we denote by $\ell_e(\cdot)$. We assume that each ℓ_e is a nonnegative, continuous, and nondecreasing function. The latency of a path P with respect to a flow f is then the sum of the latencies of the edges in the path, denoted by $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. We call the triple (G, r, ℓ) an *instance*.

Flows at Nash Equilibrium. A flow f feasible for (G, r, ℓ) is said to be *at Nash equilibrium* (or is a *Nash flow*) if for every i and every two s_i - t_i paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $\ell_{P_1}(f) \leq \ell_{P_2}(f)$. Briefly, a Nash flow routes flow on shortest paths, relative to the induced edge latencies. Also, it is well known that Nash flows always exist, and that we can assume that Nash flows are unique and are directed acyclic in single-commodity networks (see [1, 3]).

Finally, for a flow f , by $M(f)$ we mean the maximum latency $\max_{P \in \mathcal{P}: f_P > 0} \ell_P(f)$ experienced by network traffic; this is the objective function that we seek to minimize. For an instance (G, r, ℓ) admitting a flow f at Nash equilibrium, by $\rho_M(G, r, \ell)$ we mean the ratio $M(f)/M(f^*)$, where f^* is a flow feasible for (G, r, ℓ) that minimizes the maximum latency.

*Computer Science Division, UC Berkeley, Soda Hall, Berkeley CA 94720. Supported by an NSF Postdoctoral Fellowship. This research was done while the author was at Cornell University and supported by ONR grant N00014-98-1-0589. Email: timr@cs.berkeley.edu.

3 A Theorem and Two Conjectures

Single-Commodity Instances. We begin by determining the price of anarchy for the maximum latency for single-commodity networks, in which all traffic shares a common source and destination. We first prove an upper bound on the price of anarchy ρ_M , and then demonstrate that our bound is the best possible.

THEOREM 1. *If (G, r, ℓ) is a single-commodity instance with $n \geq 2$ vertices, then $\rho_M(G, r, \ell) \leq n - 1$.*

Proof. Let f be a Nash flow for (G, r, ℓ) , f^* some other flow feasible for (G, r, ℓ) , and $d(v)$ the shortest-path distance from s to v w.r.t. edge lengths $\{\ell_e(f_e)\}$. By the definition of a Nash flow, all flow paths of f are shortest s - t paths w.r.t. these edge lengths. Thus, the common (and hence maximum) latency encountered by traffic in f is precisely $d(t)$.

Since (G, r, ℓ) is a single-commodity instance, the Nash flow f can be assumed directed acyclic (see Section 2), and hence the vertices of G can be sorted in topological order w.r.t. f . We choose some such topological ordering in which $d(v)$ is nondecreasing in the ordering. Such an ordering always exists, since d -values can only increase along a sequence of edges that carry f -flow (see also [1]).

We now pick consecutive vertices v, w in the ordering that precede (or equal) t and maximize the difference $d(w) - d(v)$; since $d(t)$ is the sum of at most $n - 1$ such differences, the maximum difference is at least $d(t)/(n - 1)$. Let S be the set of vertices between s and v , inclusive, in the ordering. The set S is an s - t cut, and since vertices are sorted topologically w.r.t. f , no f -flow enters S and hence the amount of f -flow exiting S is precisely r . The s - t flow f^* must send at least r units of flow out of the cut S , so there must be an edge $e = (u, x)$ exiting S on which $f_e^* \geq f_e$ and $f_e^* > 0$, and hence $M(f^*) \geq \ell_e(f_e^*) \geq \ell_e(f_e)$. Moreover, since f sends flow on shortest paths with respect to the induced edge latencies, $\ell_e(f_e) \geq d(x) - d(u)$. Since u is or precedes v in the ordering, x is or succeeds w in the ordering, and d -values can only increase with the ordering, $d(x) - d(u) \geq d(w) - d(v)$. Thus,

$$M(f^*) \geq \ell_e(f_e) \geq d(w) - d(v) \geq d(t)/(n - 1),$$

and the proof is complete.

The bound of $n - 1$ in Theorem 1 is the best possible for all $n \geq 2$. To see this, fix $n \geq 2$, let G be the network with vertices v_1, \dots, v_n with $s = v_1$ and $t = v_n$, and with two edges, a_i and b_i , directed from v_i to v_{i+1} for each $i = 1, \dots, n - 1$. For each such i , we define $\ell_{a_i}(x) = 1$ and ℓ_{b_i} as some continuous, nondecreasing function satisfying $\ell_{b_i}(\frac{n-2}{n-1}) = 0$ and $\ell_{b_i}(1) = 1$. Then,

the flow f that routes one unit of flow on edge b_i for all i is at Nash equilibrium for $(G, 1, \ell)$ with $M(f) = n - 1$. On the other hand, the flow f^* that routes $1/(n - 1)$ units of flow on each of the $n - 1$ s - t paths of G that eschew exactly one edge of the form b_i is feasible for $(G, 1, \ell)$ with $M(f^*) = 1$.

Multicommodity Instances. While the above lower bound of $n - 1$ on the price of anarchy for the maximum latency trivially carries over to multicommodity networks, the proof of Theorem 1 relies on the combinatorial structure of a single-commodity network. We leave open the question of finding a finite upper bound on the price of anarchy for multicommodity networks, but conjecture that such a bound exists.

CONJECTURE 2. *There is a finite-valued bivariate function g such that for all multicommodity instances (G, r, ℓ) in which G has at most $n \geq 2$ vertices and $m \geq 1$ edges, $\rho_M(G, r, \ell) \leq g(n, m)$.*

We emphasize that this conjecture asserts that the price of anarchy is a function *only* of the network size, and *not* of the edge latency functions. No such bound exists for the *average* latency, even in single-commodity, two-node two-link networks [3].

We also make the stronger conjecture that Theorem 1 carries over to multicommodity networks.

CONJECTURE 3. *For all multicommodity instances (G, r, ℓ) with $n \geq 2$ vertices, $\rho_M(G, r, \ell) \leq n - 1$.*

A positive resolution of Conjecture 2 would, in particular, provide the first upper bound on the severity of “Braess’s Paradox”—the counterintuitive fact that removing edges from a network can *decrease* the latency of all traffic in a Nash flow—in multicommodity networks. The worst-case severity of Braess’s Paradox in single-commodity networks was previously determined by Roughgarden [1].

A negative resolution of Conjecture 3 would also be of interest, since it would be the first demonstration that multicommodity networks are more ill-behaved than single-commodity networks with respect to a natural price of anarchy measure. Indeed, it is known that for the price of anarchy for average latency, no such separation exists [2].

References

- [1] T. Roughgarden. Designing networks for selfish users is hard. In *FOCS '01*, pages 472–481.
- [2] T. Roughgarden. The price of anarchy is independent of the network topology. In *STOC '02*, pages 428–437.
- [3] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [4] D. Weitz. The price of anarchy. Unpublished manuscript, 2001.