

Genericity and Congestion Control in Selfish Routing

Eric J. Friedman*

School of Operations Research and Industrial Engineering,
Cornell University, Ithaca, NY 14853.

April 30, 2003

Abstract

In this paper we consider the problem of selfish routing in a congested data network, such as the Internet. While previous analyses (e.g., [RT00, Rou01]) have discussed the possibility of large losses due to selfish routing, we present several reasons why one could expect typical losses to be small. The first is based on a “generic analysis” where we consider worst case topologies and latency functions, but ignore a small set of “transmission demands.” We show that one can bound these “generic losses” by the log of the network’s “criticality factor,” a reasonably natural parameter. Our second reason is based on the near universal use of TCP or some other congestion control mechanism in networks. We show that for a specific model of TCP, the losses due to selfish routing are quite small and suggest that this is true in general. Both of these results are also shown to hold for the (nonlinear) latency functions which commonly arise on the Internet. Lastly, we provide some non-game-theoretic justifications for this analysis, which may be applicable to current networks, in which routing is not selfish. We show that in certain cases, common routing algorithms, such as OSPF, generate the same routes as selfish routing.

*I would like to thank Jon Kleinberg, Tim Roughgarden and Eva Tardos for helpful comments and Scott Shenker for suggesting the analysis in Section 4. This material is based upon work supported by the National Science Foundation under Grant No. ANI-9730162. Email: friedman@orie.cornell.edu. www: <http://orie.cornell.edu/friedman>

1 Introduction

The combination of ideas from game theory with those arising in computer science is becoming increasingly important in analyses of many problems. (See, e.g., [Nis99, NR99, FPS00, Pap01, AT01, JV01].) One of the main motivations for much of this work is the growth and privatization of the Internet. Such analyses are crucial for the Internet’s continued operation as the cooperative spirit of the Internet can not continue as it completes its transition from an academic to a business environment. For example, why should an application use a congestion control protocol when its own performance could be enhanced by disabling congestion control. (See e.g., [Nao69, San88, She90, FS98]).

In this paper we focus on the specific problem of selfish routing, in which agents get to choose the path on which their traffic will travel. Motivated by this problem [KP99] studied a scheduling model. They computed the losses due to selfish routing, which they denoted the “price of anarchy.” Recently, [MS01] and others have extended their results for related models of this type.

In “How bad is selfish routing” [RT00], Roughgarden and Tardos introduced a much more realistic model of routing on the Internet. Their model is similar to current proposals to give applications explicit control of the routes they use [CCS95].

In that paper, [RT00], they explore the efficiency losses of selfish routing. They show that the loss from selfish routing can be unbounded for general networks and latency functions. However, they do construct bounds on the losses when the latency functions are polynomials with fixed degree. They also provide an intriguing bi-criteria result which shows (essentially) that the losses due to selfish routing can be completely compensated for by doubling the bandwidth on all links.

While these are powerful results, they are somewhat unsatisfying in relation to the Internet. First, most latency functions are not polynomial. In fact they typically diverge when the traffic rate approaches the bandwidth. In this situation their analysis provides no bounds on the losses due to selfish routing. Second, it is difficult to interpret the true meaning of the bi-criteria result. One interpretation is that the losses due to selfish routing can be compensated for by doubling the bandwidth of the network. However, this is typically not a reasonable option, in the short term, and even if one could double the bandwidth then they would probably still wish to eliminate the possibly large losses due to selfish routing.

In this paper we extend Roughgarden and Tardos’ analysis in order to provide bounds which alleviate these criticisms. We consider two common aspects of data networks and the Internet, both of which lead to bounded losses.

Our first approach is to analyze a “generic version” of the problem. While we follow the standard practice of considering worst case in terms of the network latency functions and topology (for which the notion of generic or typical is highly contentious [AH01]) we do not require that the actual level of traffic be worst case and thus we consider the concept of generic¹ or typical transmission rates. Given the variability of network traffic over time, this seems like a reasonable concession for understanding the important issues involved with selfish routing.

Unfortunately, as we show, even the generic behavior in this sense can be arbitrarily bad (although the networks with this property seems quite artificial). However, if we condition our analysis on the network’s “constant of criticality” which is defined by the rate at which the Nash (selfish)

¹Note that our analysis does not apply genericity in the formal sense of considering open dense subsets of the rates, but is similar in spirit and was motivated by ideas from generic analysis in other fields (e.g., differential topology [Hir76] or game theory [FT91] ch.12).

equilibrium changes with a change in rates, then we can prove dramatically improved bounds which arise from generic analysis. In particular, if we denote this constant of criticality by C then the generic losses are bounded by $O(\log(C))$ while the losses for specific rates can be linear in C , an exponential increase. This result holds for various definitions of genericity and criticality and appears to be robust.

We also show that these results are essentially unchanged when we inject more reality into our model by restricting to latency functions which naturally arise on networks, such as the delay for an M/M/1 queue.

Our second approach considers the effect of TCP on the analysis. While Roughgarden and Tardos' model assumes that agents have a fixed transmission rate, on the Internet most flows use TCP for which the transmission rate is determined by the congestion on the network. Additionally, in the near future, even UDP flows will most likely provide congestion control under the "TCP-friendly" protocol [FHP⁺00]. Thus, we consider a version of this model under which agents do not choose a transmission rate, but instead open TCP-like connections. In this model the losses due to selfish routing are quite small.

Lastly, we note that currently on most networks, routing is typically not under control of the applications. While there have been proposals to allow versions of selfish routing on the Internet [CCS95], under the current infrastructure it is impractical. However, as we discuss more fully below, this does not make our analysis irrelevant even in the short run. Standard routing protocols with recommended metrics are actually iterative algorithms for finding the Nash equilibrium of the selfish routing game. So essentially our analysis is comparing the routing that arises on the Internet under the use of a suboptimal, but commonly used metric, to that of the optimal metric, which is much more difficult to compute. Thus, models of selfish routing might have direct implications for the current operation of the Internet.

2 Model

We consider a general network. Let $G = (V, E)$ be a graph with vertex set V , edge set E , and S a set of source-destination vertex pairs. For each $s \in S$ let $r(s)$ be the rate of flow that is required to flow between the vertices in s and let P_s be the set of simple paths between them with $P = \bigcup_s P_s$. Let $x_P \geq 0$ be the flow on path P and let $x_e = \sum_{P|e \in P} x_P$. A flow x is feasible if $\sum_{P \in P_s} x_P = r(s)$ for all $s \in S$.

When x_e is the total flow on edge e , then the latency of that link is given by $d_e(x_e)$, where $d_e(\cdot)$ is a non-negative, continuous and nondecreasing function defined on all of \mathfrak{R}_+ . The latency of a path for a flow x is given by $d_P(x) = \sum_{e \in P} d_e(x_e)$.

The cost of a flow $x = \{x_e\}$ is $\sum_{e \in E} d_e(x_e)x_e$. Then the cost of the (unique) optimal is flow is denoted $Opt(r)$ and we let $x^*(r) =$ denote the optimal flow.

However, if users (each controlling an infinitesimal amount of flow) selfishly minimize their own cost then we get the Nash equilibrium.² Let $\hat{x}(r)$ denote the Nash equilibrium flow and $N(r)$ its cost. It is straightforward to characterize this flow. In the Nash equilibrium flow we have that for all $s \in S$, $P, P' \in P_s$ such that $d_P(\hat{x}(r)) > 0$ then $d_P(\hat{x}(r)) \leq d_{P'}(\hat{x}(r))$. There is an analogous characterization for the Optimal flow.

²Formally, we are studying a non-atomic game. See e.g., [Sch73].

We will be interested in the “loss” due to selfish behavior, $\Gamma(r) = N(r)/Opt(r)$. We will prove generic bounds for this function over a set of values of r in terms of the degree of criticality of the network over this interval $C(r) = N(r)/N(r/2)$. Note, if we were only interested in a single value of r then, by a theorem of Roughgarden and Tardos [RT00], $N(r/2) \leq Opt(r)$ which is equivalent to the statement $\Gamma(r) \leq C(r)$ and it is easy to construct examples where this bound is tight. Thus, the degree of criticality directly measures the potential worst case, non-generic, losses due to selfish routing.

Although our measure of criticality seems quite natural, $C(r) = N(r)/N(r/2)$ there are other natural measures, such as $\hat{C}(r) = Opt(r)/Opt(r/2)$. However, as we discuss later, different choices for the definition of criticality do not significantly affect our results.

3 A Generic Analysis

3.1 Motivating Example

Roughgarden and Tardos consider a simple example of a parallel network with only two links. The first has constant latency $d_1(x_1) = 1$ while the second has latency $d_2(x_2) = x^k$ for some $k > 0$. They show that $\Gamma(1) \approx k/\ln(k)$ for $k \gg 1$.

However, for $r < 1$, the results are less dire. In particular, when $r < (k+1)^{-1/k} \approx 1 - \ln k/k$ the Nash flow is equal to the optimal flow and thus $\Gamma(r) = 1$. Thus as k becomes large $\Gamma(1)$ becomes large but the set of flows for which $\Gamma(r)$ is small approaches the entire interval $[0, 1]$. In this case, the flow with $r = 1$ is exceptional and for “most” values of r the loss due to selfish routing is extremely small. (Although the $N(r) \neq Opt(r)$ for any $r > 1$, the region for which $\Gamma(r)$ is large is still quite small.)

3.2 An example with generically large losses

However, in contrast to the preceding example, we now show that generic losses can be arbitrarily large. Specifically, we show that the losses due to selfish routing can be large for over the entire interval $[r/2, r]$.

Consider a parallel network. Choose some integer $n > 2$ and define $d_0(x) = 1$, $d_k(x) = \max[n^{-k}, 1 + \epsilon^{-k}(x - (1/2)/n)]$ for $0 < k \leq n$ where $\epsilon \ll 1/n$ and $d_{n+1}(x) = 1 + \epsilon^{-n-1}(x - 1/2)$. By construction, $N(1) = 1$ where $x_0 = 0$, $x_k = (1/2)/n$ for $1 \leq k \leq n$, and $x_{n+1} = 1/2$ while the optimal flow occurs when $d_k(\hat{x}_k) = n^{-k}$ for $k \leq n$ and $d_{n+1}(\hat{x}_{n+1}) = 0$ where $x_0 = O(\epsilon)$, $x_k = (1/2)/n + O(\epsilon^k)$ for all $1 \leq k \leq n$ and $x_{n+1} = 1/2 + O(\epsilon^{n+1})$. Thus, $Opt(1) < 1/(2n(n-1)) + O(\epsilon)$ and therefore $\Gamma(1) > 2n(n-1) + O(\epsilon)$. Also, $N(1/2) = n^{-n} + O(\epsilon)$ while $Opt(1/2) = O(\epsilon)$, so $\Gamma(1/2) = \Omega(1/\epsilon)$. Lastly consider the best case, when the flow is the largest value of r such that $d_1(\hat{x}_1(r)) = n^{-1}$. In this case $N(r) = n^{-1} + O(\epsilon)$, while $Opt(r) < 1/(2n(n-1)) + O(\epsilon)$ and therefore $\Gamma(r) > 2(n-1) + O(\epsilon)$. In fact, for any $r \in [1/2, 1]$, $\Gamma(r) > 2(n-1) + O(\epsilon)$.

Using this construction we can show that given any $\gamma > 0$ and $r' > r > 0$ there exists a network, such that $\Gamma(r) > \gamma$ for all $r \in [r, r']$. In particular, we can construct a network for which $\Gamma(r') \geq \gamma$ for all $r' \in [r/2, r]$. By combining multiple copies of this network we can extend this result to a general network with any number of source/destination pairs.

However, for this network note that $C(r) = N(r)/N(r/2) \approx (\gamma)^\gamma$ and thus the losses are approximately logarithmic in $C(r)$. This motivates the following analysis.

3.3 Generic Results

Our results follow from the main result in [RT00].

Proposition 1 (Roughgarden and Tardos 2000) *For all $\delta \in (0, 1]$ and r , $N(r) \leq \text{Opt}(r(1 + \delta))/\delta$.*

Our key tool will be a simple corollary of this proposition.

Corollary 1 *For all $\delta \in (0, 1]$, if $\Gamma(r) = \gamma$ then $N(r/(1 + \delta)) \leq N(r)\gamma^{-1}/\delta$.*

Our results follow from this corollary, since, it implies that for any value of r such that the losses from selfish routing are large the value of the Nash equilibria must be falling rapidly. Thus if there are a large set of points for which the losses from selfish routing are large then the cost of the Nash equilibrium must drop exponentially in the number of such points. A similar technique was used by [You98] to show that for many caching algorithms “most” choices of cache sizes lead to good performance.

As an immediate result we state the following, which shows that if $\Gamma(r)$ is large for a large set of flows, then the cost of the Nash equilibrium must grow extremely rapidly.

Theorem 1 *For any network, $r > 0$ and $\Gamma(\alpha r) > \gamma$ for all $\alpha \in [1/2, 1]$, then $\log C(r) \geq \lfloor \log(2)/\log(1 + e/\gamma) \rfloor = \gamma/e + O(1)$ and thus $C(r) = \exp[\Omega(\gamma)]$.*

Proof: Suppose that for all $r' \in \{\alpha r \mid \alpha \in [1/2, 1]\}$, $\Gamma(r') > \gamma$ and let $d = 1 + e/\gamma$. Then from the above lemma we see that for any $r' \in \{\alpha r \mid \alpha \in [1/2, 1]\}$, $N(r'/d) \leq N(r')/e$. Repeatedly applying this inequality we get $N(r/d^j) \leq N(r)/e^j$ as long as $d^j \leq e$. Since $N(r)$ is decreasing in r we get $N(r/e) \leq N(r)/e^{\lfloor \log(2)/\log(d) \rfloor}$. Plugging in the value for d yields the result. \square

The following corollary simply restates the above theorem.

Corollary 2 *For any network and $r > 0$, $\min\{\Gamma(\alpha r) \mid \alpha \in [1/e, 1]\} = O(\log C(r))$.*

We now prove our main result, that for most values of r , $\Gamma(r)$ is bounded by the logarithm of C .

Theorem 2 *For any network, $\gamma > 0$ and $r > 0$, the set of rates with losses greater than γ , $A_\gamma = \{\alpha \in [1/2, 1] \mid \Gamma(\alpha r) \geq \gamma\}$ has small Lebesgue measure,*

$$\log C(r) \geq \lfloor \nu(A_\gamma) \log(2)/\log(1 + e/\gamma) \rfloor$$

and thus

$$\nu(A_\gamma) = O(\log C(r)/\gamma),$$

where $\nu(A_\gamma)$ is the measure of set A_γ .

Proof: For any set A_γ define the sequence $\{\alpha^j\}_{j=0}^\infty$ by $\alpha^0 = \sup\{\alpha \in A_\gamma\}$ and recursively $\alpha^{j+1} = \sup\{\alpha \leq \alpha^j/d \mid \alpha \in A_\gamma\}$ if the sup exists and is 0 otherwise, where, as before, $d = 1 + e/\gamma$. Now, the sequence is defined so that $N(s^j)/N(s^{j+1}) \geq e$ when $s^{j+1} > 0$ and thus $\log(C(r)) \geq S(A_\gamma)$ where $S(A_\gamma) = \max\{j \mid s^j > 0\}$. We now show that the lowest bound on $C(r)$ occurs when A_γ is an interval of the form $[v, 1]$. Assume that $A_\gamma \neq [v, 1]$ for some $v \in [1/2, 1]$. Let s^j be as defined above.

Assume that $s^0 = 1$ if not then we can simply shift the entire set to the right and this can not increase $S(A_\gamma)$, since $(s^j, s^{j-1}/d] \cap A_\gamma = \emptyset$. Next assume that for all $j \leq S(A_\gamma)$, $s^j = s^{j-1}/d$, if not we can shift the part of A_γ contained on $[1/2, s^j]$ to the right so that this holds without increasing $S(A_\gamma)$. Last, suppose that A_γ does not include all of $[s^j, 1]$ when $j = S(A_\gamma)$. Then we can simply fill it in using some of the measure from the left side of the set, which can not increase $S(A_\gamma)$. Thus we see that intervals of the form $[v, 1]$ give the strongest bound. (Note that the sets A_γ must be finite collections of closed intervals, so all sets are measurable in the above construction.) \square

Note that by using an extension of Roughgarden and Tardos's theorem for parallel networks we can improve the above results by a factor of about 3.2. This is discussed in the Section 7.

Recall, if we are only interested in a single value of $r \in [1/2, 1]$ then the value of $\Gamma(r)$ is only bounded by $C(r)$ in contrast to our bounds which depend on $\log(C(r))$.

3.4 Other Definitions of Criticality and Genericity

Note that another natural measure of criticality would use the ratio $\hat{C}(r) = \text{Opt}(r)/\text{Opt}(r/2)$. However, this would not significantly alter our results since if $\Gamma(r) = \gamma$ then since $\text{Opt}(r/2) \leq N(r/2)$ we get $\hat{C}(r) \geq (N(r)/\gamma)/N(r/2) = C(r)/\gamma$ so $\log(\hat{C}(r)) \geq \log(C(r)) - \log(\gamma)$. Since our results show that $\gamma = O(\log C(r))$ the second term would not be significant.

Additionally, our analysis has studied the losses from selfish routing along a very particular geometric object: a portion of a ray, $[r/2, r]$. We now discuss generalizations to other geometric objects.

Firstly, note that it is trivial to extend our analysis to intervals of different lengths, such as $[\rho r, r]$ for any $\rho \in (0, 1)$. For example in Theorem 1 would apply on this interval after replacing $\log(2)$ with $\log(1/\rho)$.

Secondly, the use of rays from the origin is not crucial to our analysis. Roughgarden and Tardos's proof of Theorem 1 can be modified to show the following extension of that theorem.

Theorem 3 *For all $\delta \in (0, 1]$, $r \geq 0$ and any nonnegative vector v (of the same dimension as r) such that $|v| = |r|$, $N(r) \leq \text{Opt}(r + \delta v)/\delta$, where $|x| = \sum_k x_k$.*

Note that when $v = r$ this reduces to Roughgarden and Tardos's theorem. Now if we define $C(r, v) = N(r)/N(r - v/2)$ as long as $r - v/2 \geq 0$ then both Theorem 1 and Theorem 2 generalize in the natural way.

Lastly, these results extend easily to higher dimensional sets of rates. For example, in a network with j source/destination pairs, we could consider a generalized rectangle of rate pairs, such as $\{r' \mid r_j/2 \leq r'_j \leq r_j, \forall j\}$. In general, let S be a set of the dimension of r and define the "size" of S to be the smallest value of ρ such that there exists some r' and v such that both r' and $r' - \rho v$ are in S . Then for any set with size $1/2$, Theorem 1 would depend on the "size" of S .

Theorem 2 could similarly extended to any set with size $1/2$ if we redefined the criticality to depend on S . For example if we assume that S is the generalized rectangle, $S = \{r' \mid \forall i: r_i/2 \leq r'_i \leq r_i\}$ then Theorem 2 would apply with $C(S) = \min_{i, r'_i} N(r_i, r'_i)/N(r_i/2, r'_i)$. A wide variety of other extensions are also possible.

3.5 Improved Results for Parallel Networks

In this section we restrict our study to parallel networks. We assume that there are n links labelled 1 to n . As an aside, we note that the parallel network problem has many other interpretations. For

example, one could interpret the choice of a particular link as the choice of a website, where different sites have different congestion characteristics. Other examples include choosing a service facility [Sti92] or even choosing a nightclub [Art94]. The convergence of play to the Nash equilibrium, when there are two links, has been analyzed in this setting for a wide range of learning rules [Fri96].

Our main tool will be the use of extensions of Markov's inequality for probability distributions, if $X \geq 0$ then $Pr[X \geq \rho E[X]] \leq 1/\rho$. Consider the random variable Z which takes on the value $d_k(x_k^*(1))$ with probability $x_k^*(1)$. (By assumption $\sum_k x_k^*(1) = 1$ so this is well defined.)

We now present a simple proof of Corollary 1 for parallel networks that will form the basis for our analysis.

Alternate proof of Corollary 1 for parallel networks: Assume that $r = 1$ and $N(1) = 1$. By assumption $Opt(1) = \gamma^{-1}$ which implies that $E[Z] = \gamma^{-1}$. By Markov's lemma $Pr(Z \geq \rho\gamma^{-1}) \leq 1/\rho$. Thus, consider the flow $x'(1)$ where $x'_k(1) = x_k^*(1)$ if $d_k(x_k^*(1)) < \rho\gamma^{-1}$ and 0 otherwise. By monotonicity, it is easy to see that $N(\sum_k x'_k(1))/(\sum_k x'_k(1)) \leq \rho\gamma^{-1}$. Noting that $\sum_k x'_k(1) \leq 1 - 1/\rho$ implies that $N(1 - 1/\rho) \leq (1 - 1/\rho)\rho\gamma^{-1}$. Setting $\rho = (1 + \delta)/\delta$ completes the proof. \square

We will refine this result for parallel networks using refinements of Markov's lemma. For example, one can prove that for any $1 < \rho < \rho'$ and $a > b$ such that $a + b(1 - \rho/\rho') = 1$, then if $X \geq 0$ either $Pr[X \geq \rho E[X]] \leq a/\rho$ or $Pr[X \geq \rho' E[X]] \leq b/\rho'$. (See the appendix for details.) Using this result we can prove the following extension of Corollary 1.

Corollary 3 *For all $\delta > \delta' \in (0, 1]$ and $a > b$ such that $a + b(1 - (1 + 1/\delta)/(1 + 1/\delta')) = 1$, if $\Gamma(r) = \gamma$ then either $N(r/(1 + \delta)) \leq N(r)a\gamma^{-1}/\delta$ or $N(r/(1 + \delta')) \leq N(r)b\gamma^{-1}/\delta'$.*

We can use this to somewhat improve Theorem 1.

Theorem 4 *For any parallel network, $r > 0$ and $\Gamma(\alpha r) > \gamma$ for all $\alpha \in [1/2, 1]$, then $\log C(r) \geq \beta\gamma \log(2)/e + O(1)$, where $\beta = 1 + 2e^{-1}(1 - E^{-1}) \approx 1.5$.*

Proof: This proof parallels that of Theorem 1 but uses Corollary 3 instead of Corollary 1 with $\delta = e/\gamma$, $\delta' = e^2/\gamma$, $a = 1/\alpha$ and $b = a2e^{-1}$. The values of a and b have been chosen so that they give precisely the same bound on the reduction from $N(r)$ to $N(r/c)$ for any c . For example, each time δ is used in Corollary 3 the value of r is reduced by a factor of $(1 + e/\gamma)/a$ while the value of N is reduced by at least $1/e$, while if δ' is used then the reduction of r is $(1 + e^2/\gamma)/b$ while the reduction in N is $1/e^2$. \square

Note that we can use more complex extensions of Markov's inequality to increase this factor further. The best we have been able to find is:

Theorem 5 *For any parallel network, $r > 0$ and $\Gamma(\alpha r) > \gamma$ for all $\alpha \in [1/e, 1]$, then $\log C(r) \geq 1.2\gamma + O(1)$.*

Finally, we note that for the case of a 2 link parallel networks, an important subclass that applies to many problems arising in networks [Fri97] we get significantly stronger results. First the key lemma.

Lemma 1 *For a two link parallel network, if $\Gamma(r) \geq \gamma$ then $N(r(1 - \gamma^{-1})) \leq \gamma^{-1}N(r)$.*

Using this one gets a doubly exponential version of Theorem 1.

Theorem 6 *For any 2 link parallel network, $r > 0$ and $\Gamma(\alpha r) > \gamma$ for all $\alpha \in [1/e, 1]$, then $\log \log C(r) \geq \lfloor \log(2)/\log(1 + e/\gamma) \rfloor = \gamma/e + O(1)$ and thus $C(r) = \exp[\exp[\Omega(\gamma)]]$.*

3.6 Queuing Latencies

Perhaps the most natural latency functions to study are the well-known queuing-delay functions. The simplest and most common of these is the M/M/1 queue for which the delay is given by $d^\mu(x) = (\mu - x)^{-1}$, where μ is the service rate. It is possible that when restricted to latencies of this type that selfish behavior might not lead to unbounded losses. We now show that this is not the case since the losses can be unbounded for a large set of flow rates.

First we note that if a parallel network has k identical links then the Nash flow on each link will be the same. This is also true for the optimal flow. Thus, s copies of a link with latency $d(x)$ is effectively the same as a single link with latency $d(x/s)$. In the examples below we will allow ourselves latencies of the type $d^\mu(x/s)$ for any $s > 0$.

Consider a parallel network with one M/M/1 link with (integral) service rate $\mu \gg 1$ and $n = \mu$ links with service rate 1. Let $r = \mu - 1$. Then the Nash equilibrium is $\hat{x}_1 = 0, \hat{x}_2 = r$ and $N(r) = \mu - 1$. One can compute the optimal flow algebraically but the following estimate will be sufficient. Consider the flow where $d^\mu(x_2) = \sqrt{\mu}^{-1}$, then $x_2 = \mu - \sqrt{\mu}$ and $x_1 = \sqrt{\mu} - 1$. The value of this flow is

$$\sqrt{\mu}^{-1}(\mu - \sqrt{\mu}) + (1 - \mu^{-1/2} + \mu^{-1})^{-1}(\sqrt{\mu} - 1) = \sqrt{\mu} + O(1)$$

and thus $\Gamma(r) = \sqrt{\mu} + O(1)$. Thus if we choose μ sufficiently large we can construct an example with arbitrarily large loss.

Combining the above analysis with our earlier example with generically large losses we can construct an example using only M/M/1 latencies with generically large losses. Thus, the restriction to M/M/1 latency functions does not alter our main observations.

Note that using the same techniques we can show that for many common delay functions arising from stochastic queues we can construct a parallel network with arbitrarily large generic losses.

4 Congestion Control

Currently on most data networks the transmission rate is determined by a flow control algorithm. For example, the flow rate of a TCP connection is controlled by the congestion in the network. Recently there have been proposals to require that unicast be ‘‘TCP-friendly’’ by setting its flow rate to be comparable to TCP [FHP⁺00]. As we now show, such constraints can dramatically reduce the losses due to selfish routing. For simplicity, we will consider a parallel network with k links.

Previously we have that each infinitesimal player controls one infinitesimal unit of flow, chooses a link to send this flow on and receives the latency of that route as the ‘‘payoff’’. Then the latency of link k depends on the total mass of players, x_k which use that link.

In a model with flow control, the flow (of packets which are not dropped) on a link is a function of the mass of players which use that link. We assume that each player repeatedly opens a single TCP connection to send a file of k packets. For simplicity, we consider the static form of TCP in which the rate for a single connection is simply $\lambda = w_e/r$ where w_e is the window size for link e and r is the roundtrip time. We also assume that the roundtrip time is given by the delay formula for an M/M/1 queue with transmission rate μ_e . Then if there are x_e open TCP connections $r = 1/(\mu_e - x_e\lambda)$. Combining these formulas yields the result that $r = (x_e w_e + 1)/\mu$ and $\lambda = w\mu/(x_e w_e + 1)$.

In this setting the latency of the link is clearly not what the agent will be minimizing. It seems natural to define the ‘‘cost’’ to an agent to be the time it takes to complete a file transfer. In this

case the agent will be minimizing the cost $c_e(x_e) = k/\lambda + r = (k + w_e)/(w_e\mu_e)(x_e w_e + 1)$. Note that by varying w_e and μ_e we can construct any affine cost function $c_e(x_e) = a_e + b_e x_e$ with $a_e, b_e > 0$.

Note that the cost function in this model is the exact counterpart of the latency function in the previous model. Thus, all of the previous results on selfish routing without flow control apply to this model when we set the latency to be the cost function.

Thus, we see that the effect of flow control is to “flatten out” the cost/latency function. In this specific case, the cost/latency function is linear and thus by a result of Roughgarden and Tardos [RT00] the loss, for any flow rate and network, due to selfish routing $\gamma \leq 4/3$. Thus, the losses in a model with flow control are bounded by a small constant and perhaps the losses due to selfish routing are insignificant in this setting, when compared to other losses that might arise for other reasons. Similar results can be obtained for other latency functions, such as M/D/1 and more generally M/GI/1 functions.

5 Selfish Routing and OSPF

Note that for the standard routing protocols (which are allowed to split flows) on the Internet, such as Open Shortest Path First (OSPF), which compute shortest paths given some metric have the Nash equilibrium as the only stable outcome when the metric used for a link is simply the latency³ $d_e(x_e)$ since the condition for a Nash equilibrium is that all flow is on links with minimal cost.⁴ Whereas, it is well known that in order to compute the optimal routing these algorithms need to use the marginal cost of a link, $d(x_e d_e(x_e))/dx_e = d_e(x_e) + x_e d'_e(x_e)$ as the metric, which takes into account the interactions between links.

However, one problem with implementing the metric for optimal routing is the numerical estimate of the derivative in an extremely noisy (bursty) system, while the selfish outcome only requires estimating the mean of the delay which is significantly easier to estimate and more stable in a dynamic sense. Thus, on these grounds one might prefer to use the suboptimal metric since, as we have shown, the losses due to selfish routing might not be significant.

A Appendix: Conditional Markov Bounds

The following lemma is straightforward to prove using the analysis in [Smi95].

Lemma 2 *For any $1 < \rho_0 < \rho_1 < \dots < \rho_n$ and $x_0, x_1, \dots, x_n > 0$ such that $\sum_{j=0}^n x_j = 1$, define inductively $a_j = x_j + a_{j+1}\rho_j/\rho_{j+1}$. Then there exists the following bound:*

There exists some $j \in \{0, n\}$ such that $Pr[X \geq \rho_j E[X]] \leq a_j/\rho_j$.

From this we can prove the following useful corollary:

Theorem 7 *For any $1 < \rho_0 < \rho_1 < \dots < \rho_n$ and $a_0, a_1, \dots, a_n > 0$ such that $a_0 + \sum_{j=0}^n a_j(1 - \rho_{j-1}/\rho_j) = 1$. Then there exists the following bound:*

There exists some $j \in \{0, n\}$ such that $Pr[X \geq \rho_j E[X]] \leq a_j/\rho_j$.

³This metric is said to be “common and useful” in a CISCO FAQ. See http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/routing.htm

⁴However, one might expect convergence in this case to be problematic.

For example, one can use this to prove the following interesting bounds:

Corollary 4 *i) There exists some $j \in \{2, n\}$ such that $\Pr[X \geq 2^j E[X]] \leq [2/(n-1)]/2^j$.
ii) For any $\rho > 1$ there exists some $\rho' \in [\rho, \beta\rho]$ such that $\Pr[X \geq \rho' E[X]] \leq [1/(1+\beta)]/\rho'$.*

Proof: Part (ii) is proved most simply by taking the limit of a discrete set of evenly spaced points. \square

Theorem 5 is based on the following extension:

Corollary 5 *There exists some $p \in [.5, 10]$ such that $\Pr[X \geq \gamma e^p E[X]] \leq p e^{1-p}/(\beta' \gamma e^p)$ where $\beta' = 2e^{1/2} - (23/2)e^{1/2}/e^{10} \approx 3.30$.*

References

- [AH01] L. Adamic and B. Huberman. The web's hidden order. *Communications of the ACM*, 44(9), 2001.
- [Art94] W. B. Arthur. Inductive reasoning and bounded rationality. *American Economic Review (Papers and Proceedings)*, 84(2):406–411, 1994.
- [AT01] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, 2001.
- [CCS95] I. Castineyra, J. N. Chiappa, and M. Steenstrup. The nimrod routing architecture. Working Draft, January 1995.
- [FHP⁺00] S. Floyd, M. Handley, J. Padhye, , and J. Widmer. Equation-based congestion control for unicast applications. In *Proc. ACM Sigcomm 2000*, 2000.
- [FPS00] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. In *Proceedings of the 32nd Annual ACM Symposium on the Theory of Computing*, pages 218–227, 2000.
- [Fri96] E. J. Friedman. Dynamics and rationality in ordered externality games. *Games and Economic Behavior*, 16:65–76, 1996.
- [Fri97] E. J. Friedman. Learnability in non-atomic externality games, with applications to computer networks. (www.orie.cornell.edu/~friedman), 1997.
- [FS98] E. J. Friedman and S. Shenker. Learning and implementation in the Internet. mimeo, available from (www.orie.cornell.edu/~friedman), 1998.
- [FT91] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- [Hir76] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, 1976.
- [JV01] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 364–372, 2001.

- [KP99] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
- [MS01] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 510–519, 2001.
- [Nao69] P. Naor. On the regulation of queue size by levying tolls. *Econometrica*, 37:15–24, 1969.
- [Nis99] N. Nisan. Algorithms for selfish agents: Mechanism design for distributed computation. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 1–15, 1999.
- [NR99] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on the Theory of Computing*, pages 129–140, 1999.
- [Pap01] C. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749–753, 2001.
- [Rou01] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 472–481, 2001. Full version available at <http://www.cs.cornell.edu/timr>.
- [RT00] T. Roughgarden and E. Tardos. How bad is selfish routing? In *41st Annual Symposium on the Foundations of Computer Science*. IEEE, October 2000.
- [San88] B. Sanders. An incentive compatible flow control algorithm for rate allocation in computer networks. *IEEE Transactions on Computers*, 37:1067–72, 1988.
- [Sch73] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7:295–300, 1973.
- [She90] S. Shenker. Efficient network allocations with selfish users. In P. J. B. King, I. Mitrani, and R. J. Pooley, editors, *Performance '90*, pages 279–285. North-Holland, New York, 1990.
- [Smi95] J. Smith. Generalized chebychev inequalities: Theory and applications in decision analysis. *Operations Research*, 43:807–825, 1995.
- [Sti92] S. Stidham. Pricing and capacity decisions for a service facility: Stability and multiple local optima. *Management Science*, 38(8):1121–1139, 1992.
- [You98] N. Young. On-line file caching. In *Proc. 9th ACM-SIAM Symposium on Discrete Algorithms*, 1998.