Braess's Paradox in Large Random Graphs

Greg Valiant*
Harvard University
valiant@fas.harvard.edu

Tim Roughgarden†
Stanford University
462 Gates Building
353 Serra Mall, Stanford, CA 94305
tim@cs.stanford.edu

ABSTRACT

Braess's Paradox is the counterintuitive but well-known fact that removing edges from a network with "selfish routing" can *decrease* the latency incurred by traffic in an equilibrium flow. Despite the large amount of research motivated by Braess's Paradox since its discovery in 1968, little is known about whether it is a common real-world phenomenon, or a mere theoretical curiosity.

In this paper, we show that Braess's Paradox is likely to occur in a natural random network model. More precisely, with high probability, (as the number of vertices goes to infinity), there is a traffic rate and a set of edges whose removal improves the latency of traffic in an equilibrium flow by a constant factor. Our proof approach is robust and shows that the "global" behavior of an equilibrium flow in a large random network is similar to that in Braess's original four-node example.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

General Terms

Algorithms, Economics, Theory

Keywords

Braess's Paradox, selfish routing, random graphs

1. INTRODUCTION

Motivated largely by algorithmic issues in the Internet, there has been a tremendous surge of research activity on the interface of network optimization and economics over the past five years. Much of this recent work is motivated by the fact that well-studied economic problems become more difficult to solve once computational constraints are imposed (e.g. [11, 25]). The converse, however, holds equally true: *classical results of network optimization can fail in the presence of additional economic constraints*. There is perhaps no starker illustration of this fact than the (in)famous *Braess's Paradox* [5].

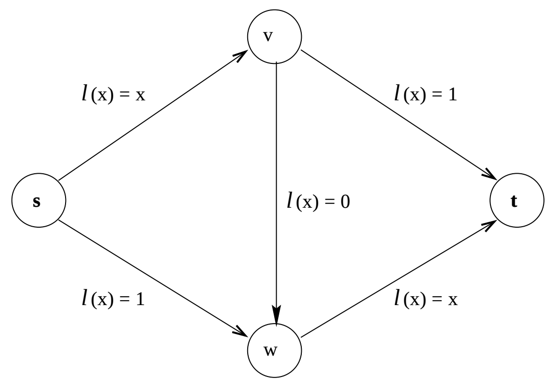

Figure 1: Braess's Paradox

1.1 Braess's Paradox

Consider the network shown in Figure 1. Assume that a large population of small network users travels from the vertex s to the vertex t, with each network user choosing an s-t path independently and selfishly (to minimize the delay experienced). Each edge of the network is labeled with its latency function, which describes the delay incurred by traffic on the link as a function of the amount of traffic that uses the link. We assume that the *traffic rate*—the total amount of traffic in the network—is 1. We also assume that traffic in the network reaches an "equilibrium flow", the natural outcome of "selfish routing" in which all traffic simultaneously travels along minimum-latency paths. In the (unique) equilibrium flow, all traffic uses the route $s \to v \to w \to t$ and experiences two units of latency. On the other hand, if we remove the edge (v, w), then in the ensuing equilibrium flow half of the traffic uses each of the routes $s \to v \to t$ and $s \to w \to t$. In this equilibrium, all network users experience latency $3/2$ and are thus better off than before.

In classical network flow, reducing the number of feasible solutions (e.g., by removing an edge of a network) trivially

*Part of this work done while visiting Stanford University and supported in part by DARPA grant W911NF-05-1-0224.
†Supported in part by ONR grant N00014-04-1-0725, DARPA grant W911NF-05-1-0224, and an NSF CAREER Award.

only degrades the optimal objective function value (for any objective function). Braess's Paradox shows that even this simple fact no longer holds for equilibrium flows in networks: removing links can improve the performance of the equilibrium flow of the network.

Since its discovery in 1968 [5], Braess's Paradox has generated an enormous amount of subsequent research in the transportation, networking, and theoretical computer science communities (see [29] and Subsection 1.3 below). However, remarkably little is known about whether Braess's Paradox is a common real-world phenomenon, or a mere theoretical curiosity. Differentiating between these two possibilities is clearly an important issue. For example, it is well known that equilibrium flows arise not only in networks with "source routing"—networks where each end user is assumed to possess knowledge of the entire network and the ability to choose an end-to-end path for its traffic—but also in networks that use a distributed delay-based routing protocol to route traffic, such as the OSPF protocol with delay as the edge metric (see e.g. [3, 15]). Largely motivated by this fact, a recent sequence of papers in the networking literature [1, 9, 19, 20, 21] has studied strategies that allocate additional capacity to a network without causing Braess's Paradox to arise—intuitively, without over-provisioning a counterproductive "cross-edge" like the edge (v, w) in Figure 1. If Braess's Paradox is a rare event in selfish routing networks, then such strategies might be largely superfluous for real-world networks. If Braess's Paradox is a widespread phenomenon, however, then the problem of adding capacity (or new edges) to a selfish routing network must be treated with care.

In summary, the following basic but poorly understood question motivates our work:

Is Braess's Paradox a "pathological" example, or a pervasive phenomenon in selfish routing networks?

1.2 Our Results

Qualitatively, our main result is the following: *in a natural random network model, Braess's Paradox occurs with high probability.* To state our results formally, define the *Braess ratio* of a network as the largest factor by which the removal of one or more edges can improve the latency of traffic in an equilibrium flow. For example, the Braess ratio of the network in Figure 1 is 4/3. For our model of random networks, we prove the following.

(R1) With high probability as $n \to \infty$, there is a choice of traffic rate such that the Braess ratio of a random network is strictly greater than 1.

Here and throughout this paper, "with high probability" means with probability tending to 1 as $n \to \infty$. Thus Braess's Paradox is in fact a fairly common occurrence in large selfish routing networks, rather than an isolated anomaly.

In fact, we prove a significantly stronger result, as follows.

(R2) There is a constant $\rho > 1$ such that, with high probability as $n \to \infty$, there is a choice of traffic rate such that the Braess ratio of a random network is at least ρ.

For each fixed number n of network vertices, the probabilities in (R1) and (R2) are with respect to the random choice of the graph and of the edge latency functions. See Subsection 2.2 for further discussion of these probability distributions. The traffic rate is not random and is chosen (as a function of n) so that it scales appropriately with the "volume" of the network. Some such scaling of the traffic rate appears to be necessary for our results; see Remark 4.4 for further discussion of this point.

The first result (R1) already answers our motivating question and indicates that Braess's Paradox is widespread: in almost all networks, for an appropriate traffic rate, there is a set of edges whose removal improves the latency of traffic in an equilibrium flow.

The second result (R2) is stronger in several respects. Most obviously, it precludes dismissing the first result on the grounds that removing edges might not significantly decrease the latency of traffic in an equilibrium flow. On the contrary, removing a set of edges can typically improve this latency by a constant factor (as $n \to \infty$). Moreover, for sufficiently simple models of random networks, we can quantify this constant. In fact, we will show that in a natural model with affine latency functions—where the worst-case Braess ratio is 4/3 [30]—a random network often has a Braess ratio that is arbitrarily close to this worst case.

Another important aspect of the second result (R2) is that proving it seems intrinsically to require an understanding of the "global" structure of a random network. In particular, our proof of (R2) will show, in a precise sense, that a random network essentially behaves like a slight generalization of the network in Figure 1. The first result (R1), by contrast, might plausibly be proved using only "local" arguments. For example, one could try to prove (R1) as follows: networks similar to that in Figure 1 occur sufficiently frequently as subnetworks in a random network, and perhaps under some additional (frequently met) conditions, removing the "cross-edge" of one or more such subnetworks improves the equilibrium flow. (It is not clear, however, that such a proof approach can be made to work; we do not know how to prove (R1) along such lines.) The second result (R2), which proves that a coordinated removal of a large fraction of a network's edges improves the latency of an equilibrium flow by a constant factor, seems impervious to arguments that do not explicitly consider the global behavior of the network.

1.3 Related Work

Several previous works have shed some understanding on the prevalence of Braess's Paradox. On the empirical side, there has been a small amount of anecdotal evidence in the transportation science literature suggesting that Braess's Paradox has occurred in certain road networks [12, 18, 24].

On the theoretical side, a number of papers have explored the ranges of parameters under which Braess's Paradox can occur; most of these, however, confined their attention to the four-node network of Figure 1 [13, 17, 26, 27] or limited generalizations [14]. Indeed, it was only recently discovered that Braess's Paradox can be more severe in large, complex networks than in Braess's original four-node example [23, 28].

Most relevant to the present work are several papers in the transportation science literature that attempt to give analytical conditions that characterize whether or not a given edge (or path of edges) is *improving*, in the sense that its removal will improve the equilibrium flow in the network.

Steinberg and Zangwill [31] and Taguchi [32] gave the earliest (independent and incomparable) results along these lines; the former paper was subsequently generalized by Dafermos and Nagurney [7]. Such analytical characterizations reduce the problem of bounding the frequency of Braess's Paradox in a random network model to the (possibly easier) problem of bounding the likelihood that a certain analytical condition holds. This potential application was explicitly pointed out by Steinberg and Zangwill [31], who also noted that the form of their analytical characterization of improving edges suggested that Braess's Paradox should be common rather than rare.

The approach of [7, 31, 32] suffers from several drawbacks, however. First, the ambitious goal of analytically characterizing improving edges led to strong extra hypotheses in all of these papers. In particular, the analyses in [7, 31, 32] all assume the following when characterizing whether or not an edge e is improving: removing the edge e does not cause new s-t paths to carry traffic. Put differently, the assumption is that while removing the edge might increase the amount of traffic on other paths, it should not fundamentally change the traffic pattern. This assumption clearly fails e.g. in the network of Figure 1, and it is not clear that it typically holds in large random networks. It is singled out by Steinberg and Zangwill [31, §7] as the key open issue in their analysis.

Second, even when this additional hypothesis holds, it is not clear that analyzing the probability that the (somewhat complex) conditions of [7, 31, 32] hold is more tractable than directly analyzing the probability that Braess's Paradox occurs. While the condition of Steinberg and Zangwill [31] suggests that this probability could be large under the above hypothesis—it essentially states that whether or not a given edge is improving is governed by the parity of a seemingly unrelated combinatorial quantity—rigorously analyzing this condition in random graphs does not appear to be easy. In particular, at no point do Steinberg and Zangwill [31] explicitly define a probability measure over networks and analyze the probability that a given edge is improving.

Third, all of the above characterizations consider only the effects of removing a single edge of a network. Even under strong additional assumptions, this "local" approach seems incapable of proving an analogue of our second result (R2), which shows that the coordinated deletion of a large set of edges yields a constant-factor improvement in equilibrium flow latency.

In summary, we believe the present paper to be the first to explicitly define a natural probability distribution over selfish routing networks and analyze the probability that Braess's Paradox occurs, and the first to consider the simultaneous deletion of multiple edges or to quantify the Braess ratio in large random networks.

2. THE MODEL

2.1 Selfish Routing Networks

We follow the notation and conventions of Roughgarden and Tardos [30]. We study a single-commodity flow network, described by a graph $G = (V, E)$ with a source vertex s and a sink vertex t. We assume for convenience that all graphs are undirected, although allowing directed graphs would not affect our results in any significant way. We denote the set of simple s-t paths by $\mathcal{P}$, and we assume that this set is nonempty. A *flow* f is a nonnegative vector, indexed by $\mathcal{P}$. For a fixed flow f we define $f_e = \sum_{P\in\mathcal{P}:e\in P} f_P$ as the amount of traffic using edge e en route from s to t. With respect to a finite and positive *traffic rate* r, a flow f is said to be *feasible* if $\sum_{P\in\mathcal{P}} f_P = r$.

We model congestion in the network by assigning each edge e a nonnegative, continuous, nondecreasing *latency function* ℓ_e that describes the delay incurred by traffic on e as a function of the edge congestion f_e. The latency of a path P in G with respect to a flow f is then given by $\ell_P(f) = \sum_{e\in P} \ell_e(f_e)$. We call a triple (G, r, ℓ) an *instance*.

In Section 1 we informally discussed equilibrium flows; we now make this notion precise.

Definition 2.1 ([33]) A flow f feasible for (G, r, ℓ) is *at Nash equilibrium* or is a *Nash flow* if for all $P_1, P_2 \in \mathcal{P}$ with $f_{P_1} > 0$, $\ell_{P_1}(f) \le \ell_{P_2}(f)$.

Thus all paths in use by a flow at Nash equilibrium have equal latency. As is well known, every selfish routing network admits at least one Nash flow [2]. Moreover, Nash flows are "essentially unique" in the sense that the latency incurred by traffic is the same in every Nash flow of a network [2]. We use the notation $L(G, r, \ell)$ to denote the common latency of all traffic in a flow at Nash equilibrium for the instance (G, r, ℓ).

The following (well-known) characterization of Nash flows will be instrumental in our proofs. It follows easily from the fact that a flow at Nash equilibrium routes traffic only on minimum-latency paths.

Proposition 2.2 ([28]) *Let f be a flow feasible for (G, r, ℓ). For a vertex v in G, let $d(v)$ denote the length, with respect to edge lengths $\ell_e(f_e)$, of a shortest s-v path in G. Then $d(w) - d(v) \le \ell_e(f_e)$ for all edges $e = (v, w)$, and f is at Nash equilibrium if and only if equality holds whenever $f_e > 0$.*

We will also use the intuitive but non-obvious fact that the latency $L(G, r, \ell)$ of traffic in a Nash flow is continuous and non-decreasing in the traffic rate r.

Proposition 2.3 ([16, 22]) *For every fixed network G and latency functions ℓ, the value $L(G, r, \ell)$ is continuous and nondecreasing in r.*

2.2 Models of Random Networks

In order to rigorously claim that Braess's Paradox is or is not likely to occur, we must fix a model of random selfish routing networks. Such a model contains (at least) two ingredients: a probability distribution over graphs and a probability distribution over edge latency functions. While the field of random graph theory (e.g. [4]) provides a vast array of possible definitions of and analytical tools for random graphs, choices for the definition of a "random latency function" are less obvious. In this paper, we make the following two basic modeling assumptions.

(1) The underlying graph G is distributed according to the standard Erdös-Renyi $\mathcal{G}(n, p)$ model [10]. Precisely, for a fixed number $n \ge 2$ of vertices, we assume that each possible (undirected) edge is present independently with probability p. We also assume that $p = \Omega(n^{-1/2+\epsilon})$ for some $\epsilon > 0$. The source s and the sink t are chosen randomly or arbitrarily. Finally,

to avoid degenerate cases, we assume that there is no direct (s, t) edge.

(2) Latency functions are *affine*—of the form $\ell(x) = ax+b$ with $a, b \geq 0$. (Such functions are often called *linear* latency functions.)

We make the first assumption simply because the Erdös-Renyi model is the most popular and widely studied definition of a random graph. Our proof techniques do not crucially use detailed properties of this model, however, and we suspect that they are general enough to apply to every random graph model where a typical graph is "sufficiently dense and uniform". We defer rigorous analysis of this suspicion to future work. Whether or not our results carry over to models of sparse or non-uniform random graphs is an interesting open question.

Our motivation for assumption (2) is that affine latency functions are, informally, the most benign functions that allow Braess's Paradox to occur. More precisely, in networks with only constant latency functions or with only affine latency functions with zero constant terms, deleting edges can only increase the cost of a flow at Nash equilibrium [8]. On the other hand, allowing nonlinear latency functions only increases the worst-case severity of Braess's Paradox [28]. For example, the network of Figure 1 has the largest-possible Braess ratio among all networks with affine latency functions [30], but larger Braess ratios are possible in networks with nonlinear latency functions [28].

Since our goal is to lower bound both the frequency and severity of Braess's Paradox in random networks, our restriction to the relatively benign class of affine latency functions is well motivated. Moreover, it will be intuitively clear that our analysis approach is robust enough to extend, with some work, to natural models of random nonlinear latency functions.

Even for affine latency functions, there are many possible definitions of a random latency function. We will focus most of our attention (Section 4) on the *independent coefficients* model. In this model, we assume that there are two fixed distributions $\mathcal{A}$ and $\mathcal{B}$, and each edge is independently given a latency function $\ell(x) = ax + b$, where a and b are drawn independently from $\mathcal{A}$ and $\mathcal{B}$, respectively. We prove our main result for this model—that with high probability, removing some set of edges improves the latency of a Nash flow of a random network by a constant factor—under mild assumptions on the distributions $\mathcal{A}$ and $\mathcal{B}$.

We also consider what we call the $1/x$ *model*, where each edge present in the graph (independently) has the latency function $\ell(x) = x$ with probability q and the latency function $\ell(x) = 1$ with probability $1 - q$. Note that this model is not a special case of the independent coefficients model, since there is now (complete) dependence between the a- and b-coefficients of the latency function of an edge. While stylized, this model nevertheless serves several purposes: it shows that independence of coefficients is not essential for our earlier results; it provides a clean example of how our high-level proof approach can be adapted to different random network models; and we can obtain a precise bound on the Braess ratio of a random network in this model (as a function of the parameters p and q). In particular, for sufficiently small p and q, we prove that a random network in this model is essentially a worst-possible example of Braess's Paradox (among networks with affine latency functions).

3. HIGH-LEVEL PROOF APPROACH

In this section we describe our high-level proof approach. Assume for simplicity that the random graph parameter p is bounded below by a constant; this is for the sake of intuition only and is not required for our results.

At the highest level, our plan is to show that a random network has a "global" structure similar to that of the four-node network of Figure 1. To make this more precise, recall the distance labels of Proposition 2.2: for an instance (G, r, ℓ), let $d(v)$ denote the length of a shortest s-v path with respect to the edge latencies induced by a Nash flow of (G, r, ℓ). In the network of Figure 1, we have $d(s) = 0$, $d(v) = d(w) = 1$, and $d(t) = 2$. After removing the edge (v, w), the distance labels become $d(s) = 0$, $d(v) = 1/2$, $d(w) = 1$, and $d(t) = 3/2$.

Now consider a large random graph G, under some random network model. We will choose the traffic rate to scale appropriately with the size of G and consider a flow at Nash equilibrium in G. Ignore vertices unused by this Nash flow and label the other vertices $v_1, \ldots, v_k$ so that $d(v_1) \leq \cdots \leq d(v_k)$. Proposition 2.2 implies that, without loss of generality, $s = v_1$ and $t = v_k$. A key step in our analysis, which we call the "Delta Lemma", is to show that $d(v_2) \approx d(v_{k-1})$, in the sense that $d(v_{k-1}) - d(v_2) \ll d(v_2)$, with high probability. In other words, all "internal vertices" $v_2, \ldots, v_{k-1}$ have relatively equal distance from the source (and the sink). Intuitively, the Delta Lemma holds because there are a quadratic number of "internal" edges (edges with endpoints v_i, v_j, $2 \leq i, j, \leq k-1$) but only a linear number of edges incident to the source and sink. We can thus regard G as essentially two sets of parallel links with a small latency of $\delta = d(v_{k-1}) - d(v_2)$ associated to the center node (with respect to the flow at Nash equilibrium). See Figure 2.

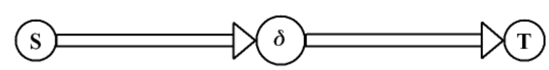

Figure 2: The "Delta Lemma". With high probability, a random network essentially behaves like two sets of parallel links with a small latency of δ associated to the center node (with respect to a flow at Nash equilibrium).

Next, for each of the two sets of parallel links, partition the links into the following three groups. First are the edges with a latency function with a constant term (b-coefficient) that is at least a parameter $\beta_1 > d(v_2) + \delta$ and an a-coefficient that is at most a parameter α_1; by Proposition 2.2 and the definition of δ, these edges carry no flow in the Nash flow of G. Second are the edges with a latency function with a constant term that is at most a constant β_2 that is significantly smaller than $d(v_2)$ and an a-coefficient that is at least a parameter α_2. Edges in these two groups are analogous to the edges with the latency functions $\ell(x) = 1$ and $\ell(x) = x$ in Figure 1, respectively. Third are the remaining edges. Figure 3 shows the network G following this partitioning; edges not relevant to our proof are omitted.

The next step, which is the most delicate step in the proof, is to define an appropriate subnetwork G' in which the Nash flow has smaller latency than in the original network G. Intuitively, we obtain G' by deleting edges from G in order to pair up the (unused) edges with latency function roughly

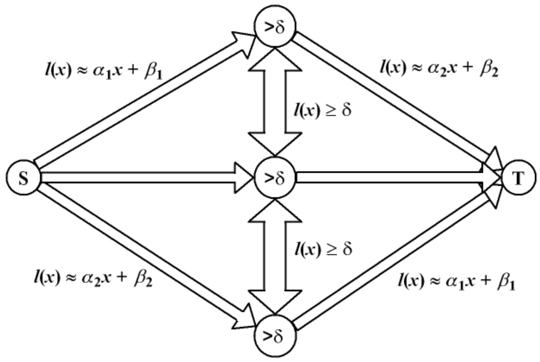

Figure 3: The network G following the Delta Lemma and the partitioning of edges according to the coefficients of their latency functions. Edges that are not relevant for our arguments are omitted.

$\alpha_1 x + \beta_1$ with those with latency function roughly $\alpha_2 x + \beta_2$. These edge deletions are analogous to the removal of the internal edge in Figure 1; see Figure 4. While the obvious hope is that removing these edges will result in a network with an improved flow at Nash equilibrium, this is not intuitively clear. On the one hand, traffic is now distributed over more paths, as in the improvement for the four-node network of Figure 1. On the other hand, we have accomplished this by selectively pairing previously unused edges with previously used edges, which in general destroys some s-t paths that contained only edges with small b-coefficients and that carried a significant amount of traffic. Proving that the edge removals result in a superior Nash flow thus requires a careful analysis that compares the benefit of employing a larger number of flow paths with the cost that many of these paths may contain edges with larger b-coefficients than those of the edges used by the Nash flow in the original network.

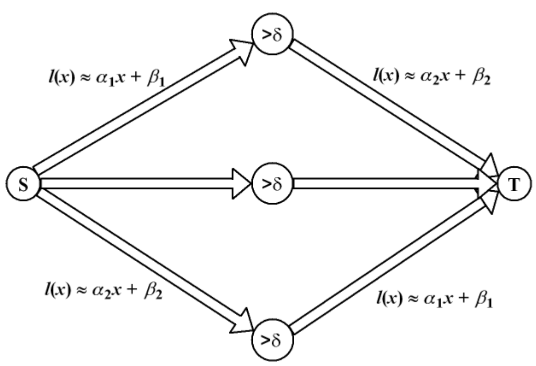

Figure 4: The network G' obtained from G by deleting all "cross edges".

In Sections 4 and 5, we prove that the above intuitive explanation for the presence of Braess's Paradox can be formalized to give constant factor expected improvements in natural random network models.

In summary, our high-level proof approach decomposes into the following four main steps (after the choice of an appropriate traffic rate).

(1) We identify a set of "good properties" and show that our networks satisfy these properties with high probability. These properties ensure that we are not dealing with anomalous instances of our random network model.

(2) [Delta Lemma] In the notation above, we prove that $d(v_2) \approx d(v_{k-1})$, and thus we may regard the entire interior section of our network as a single node with some small latency δ associated with it. We also prove an analogous result for the network G' (obtained from G via edge removals as above).

(3) [Balance Lemma] We will prove that the Nash flows in G and G' are relatively symmetric. Specifically, we show that $d(v_2) \approx d(v_k) - d(v_{k-1})$. Viewing our network as in Figure 2, the latency of the Nash flow is thus equally balanced between the left and right halves.

(4) Finally, we evaluate the latency of traffic in Nash flows in G and G', and show that the increased number of flow paths more than compensates for the decrease in flow along edges that were used in G and are now paired with costly edges in G'.

4. THE INDEPENDENT COEFFICIENTS MODEL

In this section, we apply the proof approach of Section 3 to the independent coefficients model of Subsection 2.2. After discussing preliminaries in Subsection 4.1, Subsections 4.2–4.5 formalize the four steps of this proof approach in turn.

4.1 Preliminaries

As discussed in Subsection 2.2, we assume that the underlying graph G is drawn from $\mathcal{G}(n,p)$ with $p = \Omega(n^{-1/2+\epsilon})$ for some $\epsilon > 0$ and that each edge latency function has the form $\ell(x) = ax + b$ where a and b are drawn (independently) from distributions $\mathcal{A}$ and $\mathcal{B}$. We impose some mild restrictions on $\mathcal{A}$ and $\mathcal{B}$, as follows.

Let a and b be random variables chosen from distributions $\mathcal{A}$ and $\mathcal{B}$, respectively. The distributions $\mathcal{A}, \mathcal{B}$ are *reasonable* if: (1) there is a constant $A_0 > 0$ such that $Pr(a < A_0) = 0$; (2) there is a constant $A_{\max} > 0$ such that $Pr(a > A_{\max}) = 0$; (3) there exists some interval, of length $L_A > 0$, with left endpoint A_L such that for any $\epsilon > 0$, and any $L < L_A$, $Pr(A_L + L < a < A_L + L + \epsilon) > 0$; and (4) there is a constant $L_B > 0$ such that for all $\epsilon > 0$, and all L, $0 \le L < L_B$, $Pr(L < b < L + \epsilon) > 0$.

The first assumption states that a-coefficients should be bounded away from 0. This assumption is necessary as without it, a random network is likely to contain an s-t path with essentially zero latency. In this case, Braess's Paradox will not occur. The second assumption states that a-coefficients should be bounded above by some constant; the third states that a-coefficients should be at least somewhat dense in some finite interval. All of these technical assumptions are quite weak. The last assumption states that b-coefficients should be somewhat dense near 0; while this is a stronger assumption than the previous three, it is still satisfied by most natural continuous distributions. In Section 6 we discuss the extent to which this assumption can be relaxed.

Our main result for the independent coefficients model is the following.

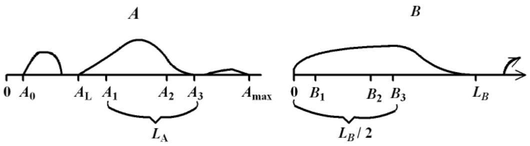

Figure 5: The parameters associated with the distributions A and B.

Theorem 4.1 *Let $\mathcal{A}$ and $\mathcal{B}$ be reasonable distributions. There is a constant $\rho = \rho(\mathcal{A}, \mathcal{B}) > 1$ such that, with high probability, a random network (G, ℓ) admits a choice of traffic rate r such that the Braess ratio of the instance (G, r, ℓ) is at least ρ.*

4.2 Properties of G

We now implement the first step of the proof approach and state several properties of a random network that hold with high probability. Throughout this subsection, we fix reasonable distributions $\mathcal{A}$ and $\mathcal{B}$ and a sufficiently large value of n.

We begin by introducing some additional parameters. Recall the meaning of the constants A_0, $A_{\max}$, L_A, A_L, and L_B from the definition of reasonable distributions. Let $\epsilon_A \ll L_A$ be a sufficiently small constant. Define $A_1 = A_L + \epsilon_A$, $A_2 = A_L + L_A - \epsilon_A$, and $A_3 = A_L + L_1$. Define $B_3 = L_B/2$ and choose $B_1 > 0$ smaller than $[B_3(L_A - 3\epsilon_A)]/[2(A_1 + A_3) + 3(L_A - 3\epsilon_A)]$. Let B_2 denote $B_3 - B_1$.

For fixed $\tau, \gamma > 0$, we now list four properties that might or might not be satisfied by a random network G. All of these properties essentially state that a simple random variable takes on a value reasonably close to its expectation.

(P1) There are at least $(np \cdot \Pr[b < B_2/3])/8$ edge-disjoint 3-hop paths between s and t that comprise only edges with b-coefficients less that $B_2/3$.

(P2) For every pair of nodes v_1, v_2, at least $(np^2 \cdot \Pr[b < \gamma]^2)/2$ other nodes w are neighbors of both v_1 and v_2, where the edges (v_1, w) and (v_2, w) each possess a latency function with b-coefficient at most γ.

(P3) For nonnegative integers i and j, let $p_{i,j}$ denote the probability that the a- and b-coefficients of a random latency function lie in the intervals $I_i = [i\tau, (i+1)\tau]$ and $I_j = [j\tau, (j+1)\tau]$, respectively. Then for all pairs i, j with $j\tau \le 2B_2$, and for each $w \in \{s, t\}$, the number of edges incident to w with a- and b-coefficients in I_i and I_j, respectively, lies in $[p_{i,j}(np - (np)^{2/3}), p_{i,j}(np + (np)^{2/3})]$.

(P4) For each $w \in \{s, t\}$, the number of edges incident to w with b-coefficient less than B_3 is at least $np(1 - (np)^{-1/3}) \cdot \Pr[b < B_3]/2$. Moreover, the total number of edges incident to w is at most $np(1 + (np)^{-1/3})$.

To define the subnetwork G' of G, we group the vertices of G into sets according to the a- and b-coefficients on the edges connecting them to the source and sink (if any). Toward this end, for a node $v \in V \backslash \{s, t\}$, let $a_s(v)x + b_s(v)$ and $a_t(v)x + b_t(v)$ denote the latency functions of the edges (s, v) and (v, t), respectively. (We will usually suppress the dependence on v in our notation.) If one or both of these edges are absent from G, we define the corresponding coefficients to be $+\infty$.

Now group vertices as follows. First assign vertices with $a_s < A_1$ and $b_s \in (B_2, B_3)$ to the set S_1, and with $a_t < A_1$ and $b_t \in (B_2, B_3)$ to the set T_1. Vertices with $a_s \in (A_2, A_3)$ and $b_s < B_1$ are assigned to the set S_2, and those with $a_t \in (A_2, A_3)$ and $b_t < B_1$ are assigned to the set T_2. Vertices in $S_1 \cap T_1$ or $S_2 \cap T_2$ are removed from both sets and placed in a "catch-all" set U. We also evacuate vertices from the largest three sets among S_1, S_2, T_1, T_2 to U until the four sets have equal size.

Vertices v with $b_s, b_t > B_3$ are distributed evenly among three sets Q_1, Q_2, Q_3, using (for example) a predetermined lexicographic rule. Finally, all remaining nodes other than s and t are placed in U. Note that, by construction, all of these sets are pairwise disjoint except possibly for the pairs S_1, T_2 and S_2, T_1.

Obtain the subnetwork G' or G by retaining only the edges whose endpoints satisfy one of the following conditions:

1. one is the source or sink, the other is not in Q_i (any $i = 1, 2, 3$);
2. each is in a different set from among S_1, T_2, or Q_1;
3. each is in a different set from among S_2, T_1, or Q_2;
4. both are in $Q_3 \cup U$.

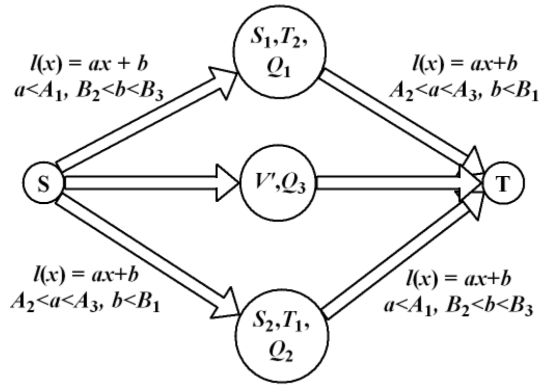

Figure 6: Depiction of the subnetwork G'.

The resulting subnetwork G' is depicted in Figure 6. We now state our remaining three desired properties. Recall that $\gamma, \tau > 0$ are arbitrary fixed constants.

(P5) Let $q_1 = \Pr[a < A_1, B_2 < b < B_3]$, $q_2 = \Pr[A_2 < a < A_3, b < B_1]$, and $q = \min\{q_1 - q_1^2, q_2 - q_2^2\}$. Then the common size of S_1, S_2, T_1, T_2 lies in $[npq/2, np]$. Additionally, the common size of each Q_i is at least $(n \cdot \Pr[b > B_3]^2)/6$.

(P6) Let v_1, v_2 be a pair of nodes that both lie in $S_1 \cup T_2$, in $S_2 \cup T_1$, or in $U \cup \{s, t\}$. Then the number of other nodes w of Q_1, of Q_2, or of Q_3, respectively, for which the edges (v_1, w) and (v_2, w) both possess a latency function with b-coefficient less than γ is at least $(np^2 \cdot \Pr[b > B_3]^2 \Pr[b < \gamma]^2)/12$.

(P7) Define intervals as in property (P3). For nonnegative integers i, j, let $h_{i,j,s}$ and $h_{i,j,t}$ denote the number of edges incident to s or t, respectively, that are also incident to a vertex of U, and that have a- and b-coefficients in the intervals I_i and I_j, respectively. Then for all i, j, $|h_{i,j,s} - h_{i,j,t}| \le (pn)^{2/3}$.

We call a network satisfying (P1)–(P7) *good*.

Lemma 4.2 *For reasonable distributions A, B and fixed constants $\gamma, \tau > 0$, a random n-node graph G is good with probability approaching 1 as $n \to \infty$.*

The proof of Lemma 4.2 is a relatively straightforward application of Chernoff bounds, and we omit the details from this extended abstract.

Recall from Definition 2.1 that all flow paths of a Nash flow in an instance (G, r, ℓ) have equal latency $L(G, r, \ell)$. Given a good network (G, ℓ), we will choose the traffic rate R so that $L(G, R, \ell) = 2B_2$. Since a-coefficients are bounded away from zero, Proposition 2.3 implies that such a traffic rate must exist. The next lemma shows that, even though R is defined implicitly, we can accurately predict its magnitude.

Lemma 4.3 *Let (G, ℓ) be a good network and let R be a traffic rate such that $L(G, R, \ell) = 2B_2$. Then,*

$$\frac{B_2 np \Pr[b < B_2/3]}{48A_{\max}} \le R \le \frac{2B_2 np(1 + (np)^{-1/3})}{A_0}. \qquad (1)$$

Proof. The upper bound follows easily from the second part of property (P4) and the definition of A_0. For the lower bound, recall that property (P1) states that there are at least $\kappa = (np \cdot \Pr[b < B_2/3])/8$ edge-disjoint 3-hop paths between s and t that comprise only edges with b-coefficients less that $B_2/3$. Splitting r units of traffic evenly between these paths yields a flow in which all traffic experiences at most $3(A_{max}r/\kappa + B_2/3)$ latency. Since the price of anarchy in networks with linear latency functions is 4/3 [30] and G is a single-commodity network, the common latency $L(G, r, \ell)$ of a Nash flow in (G, r, ℓ) is at most $4(A_{max}r/\kappa + B_2/3)$. Since this is strictly less than $2B_2$ when $r < B_2\kappa/6A_{max}$, the proof is complete. □

Remark 4.4 As we have previously noted, we assume that the graph and edge latency functions are random while the traffic rate is adversarially chosen. One could also consider a random traffic rate (according to some distribution), but there is healthy evidence that Braess's Paradox is unlikely to occur across a wide range of traffic rates (see [15, 26]). It therefore seems essential for any result similar to Theorem 4.1 that the traffic rate is carefully chosen. On the other hand, our proof does allow for a certain amount of variability in the traffic rate (to a degree depending on the distributions A, B).

4.3 The Delta Lemma

Next we prove the Delta Lemma, which states that the "internal nodes" of a good network are all of roughly the same distance from the source s.

Lemma 4.5 (Delta Lemma) *Let $\gamma > 0$ be a fixed constant and (G, ℓ) a sufficiently large good network. Define R as in Lemma 4.3 and let f be a Nash flow for (G, R, ℓ). Define $s = v_1, v_2, \ldots, v_k = t$ as in Section 3. Then*

$$d(v_{k-1}) - d(v_2) \le \frac{16A_{\max}B_2}{A_0 np^2(\Pr[b < \gamma])^2} + 2\gamma$$

Proof. First, since no edges have zero latency, Proposition 2.2 implies that all flow of f that enters v_2 and exits v_{k-1} arrives directly from s and departs directly for t, respectively. Thus the total amount of flow through either v_2 or v_{k-1} is at most $2B_2/A_0$. Next, let κ denote the number of two-hop v_2-v_{k-1} paths whose edges both have b-coefficients at most γ. By property (P2) of good networks, $\kappa \ge (np^2 \cdot \Pr[b < \gamma]^2)/2$. By an averaging argument, both edges on one of these paths each carry at most $4B_2/\kappa A_0$ units of flow. The latency of each of these edges (with respect to f) is therefore at most $\gamma + 4B_2A_{max}/\kappa A_0$; Proposition 2.2 then implies that $d(v_{k-1}) - d(v_2) \le 2\gamma + 8B_2A_{max}/\kappa A_0$, proving the lemma. □

Since $p = \Omega(n^{-(1/2)+\epsilon})$, Lemma 4.5 can be coarsely summarized as: for every pair $v, w \in \{v_2, \ldots, v_{k-1}\}$, $|d(w) - d(v)| \le 2\gamma + o(1)$. A similar argument using properties (P5) and (P6) of good networks (that we omit) proves the following analogue of the Delta Lemma for the subnetwork G'.

Lemma 4.6 (Delta Lemma for G') *Let $\gamma > 0$ be a fixed constant and (G, ℓ) a sufficiently large good network. Define R as in Lemma 4.3 and the subnetwork G' as in Subsection 4.2. Let f be a Nash flow for (G', R, ℓ), and let $d'(v)$ denote the length (w.r.t. edge lengths $\ell_e(f_e)$) of a shortest s-v path in G'. Let v, w denote two vertices that both lie in $S_1 \cup T_2$, in $S_2 \cup T_1$, or in U. Then $|d'(v) - d'(w)| \le 2\gamma + o(1)$.*

4.4 The Balance Lemma

Next we prove the Balance Lemma, which states that the latency along Nash flow paths in G is equally split between the two "halves" of the network.

Lemma 4.7 (Balance Lemma) *Let $\gamma, \tau > 0$ be fixed constants and (G, ℓ) a sufficiently large good network. Define R as in Lemma 4.3 and let f be a Nash flow for (G, R, ℓ). Define $s = v_1, v_2, \ldots, v_k = t$ as in Section 3. Let $\delta = d(v_{k-1}) - d(v_2)$. Then for all $i \in \{2, \ldots, k-1\}$,*

$$|B_2 - d(v_i)| \le \frac{4(np)^{2/3}B_2}{N} + \frac{\tau(B_2 + \delta)}{2A_0} + \delta, \qquad (2)$$

where $N = C(np - (np)^{2/3})$ and $C = C(\mathcal{A}, \mathcal{B})$ is a constant.

The proof of the Balance Lemma requires the following lemma, which argues that all edges incident to s or t with sufficiently small b-coefficients carry traffic in a Nash flow.

Lemma 4.8 *With the same assumptions and notation as Lemma 4.7, let $d(v_2) = B_2 - \sigma$ and assume that $\sigma > 2(\delta + \gamma)$. Then every edge incident to s with b-coefficient less than $B_2 - \sigma - \delta - \gamma$ carries traffic in the Nash flow f, and every edge incident to t with b-coefficient less than $B_2 + \sigma - \delta - \gamma$ carries traffic in f.*

Proof. We prove only the first assertion; the second follows from a similar argument. Let (s, v) be an edge with b-coefficient less than $B_2 - \sigma - \delta - \gamma$. By Property (P2) of good networks, there is a path $P = s \to v \to w \to t$ in which the second two edges have b-coefficients at most γ.

First, observe that edge (v,w) has latency at most γ if it carries no flow. By Proposition 2.2 and the definition of δ, it has latency at most δ if it does carry flow. Next, recall that the traffic rate R is chosen so that $L(G,R,\ell) = d(t) = 2B_2$. Proposition 2.2 and the definition of σ imply that when the edge (w,t) carries flow, it must have latency at most $B_2+\sigma$. If edge (w,t) carries no flow, then it has latency at most $\gamma < B_2+\sigma$. In all cases, the combined latency of the edges (v,w) and (w,t) (with respect to f) is at most $B_2+\sigma+\delta+\gamma$. Since the combined latency of the path P is at least $2B_2$ with respect to f and the b-coefficient of edge (s,v) is less than $B_2-\sigma-\delta-\gamma$, it must carry traffic in f. □

We now prove the Balance Lemma.

Proof of Lemma 4.7: Since $d(v_2) \leq \cdots \leq d(v_{k-1})$, it suffices to show that $B_2 - d(v_2)$ and $d(v_{k-1}) - B_2$ are both at most the right-hand side of (2). We only prove the former inequality; the other case is symmetric (using an obvious variant on Lemma 4.8).

Let R_{min} denote the left-hand side of (1). Let κ denote the number of edges incident to s; $\kappa \leq np(1+(np)^{-1/3})$ by Property (P4). Since all a-coefficients are at least A_0, the minimum cost of a flow-carrying edge of G is at least $A_0 R_{min}/\kappa$. For n sufficiently large, this is greater than the constant $c = A_0 B_2 \Pr[b < B_2/3]/96A_{\max}$. Assume that τ is a sufficiently small constant, less than $c/2$.

For n sufficiently large, property (P3) guarantees that there are at least $N := \Pr[b < c/2](np - (np)^{2/3})$ edges incident to the source with b-coefficient less than $c/2$. From Lemma 4.8, for every $\delta, \gamma < c/4$, all of these edge carry flow. By a similar argument, there are also be at least N edges carrying flow to the sink.

Define σ by $d(v_2) = B_2 - \sigma$. By the Delta Lemma (Lemma 4.5), every flow-carrying edge out of s has latency at most $B_2 - \sigma + \delta$, while every flow-carrying edge into t has latency at least $B_2 + \sigma - \delta$.

Define the interval I_i to be $[i\tau, (i+1)\tau]$. We say that an edge e has *type* (i,j) if the a- and b-coefficients of its latency functions lie in I_i and I_j, respectively. A type-(i,j) flow-carrying edge out of s carries at most $(B_2 - \sigma + \delta - j\tau)/i\tau$ flow, while a type-(i,j) flow-carrying edge into t carries at least $(B_2+\sigma-\delta-(j+1)\tau)/(i+1)\tau$ flow. By Property (P3), for all types (i,j), there are nearly the same number of type-(i,j) edges incident to each of s and t. As argued above, all such edges carry flow when the interval I_j contains only b-coefficients that are at most $c/2$.

Calculating the difference between the upper bound of flow leaving the source and the lower bound of flow entering the sink, requiring that this difference be nonnegative, and some algebra yields the inequality

$$\frac{N\tau(B_2+\delta)}{A_0^2} + \frac{2N(\delta-\sigma)}{A_0} + F_1 + F_2 \geq 0,$$

where F_1 is the error term corresponding to the variance in property (P3), and F_2 corresponds to flow using edges incident to the source with a b-coefficient that is at least $B_2 - \sigma - \delta - \gamma - \tau$.

The term F_1 is bounded above by $4(np)^{2/3}B_2/A_0$, since every edge incident to s has latency at most $2B_2$. The term F_2 is also bounded by this quantity, because all edges incident to the sink with b-coefficients less than $B_2+\sigma-\delta-\gamma$ carry positive flow, and this is at least $B_2-\sigma+\delta+\tau$, which is the maximum value of a b-coefficient of a flow-carrying edge out of the source.

Plugging these values into the inequality, and solving for σ shows the desired result. □

The following similar result holds for the subnetwork G' (formal statement and proof omitted). Define the shortest-path distance $d'(v)$ of v in G' as in Lemma 4.6. Let v_2 and v_{k-1} be the vertices of U with minimum and maximum d'-values, respectively. Then, choosing τ sufficiently small, for n sufficiently large, and assuming that (G,ℓ) is a good network, both $L(G',R,\ell)/2 - d'(v_2)$ and $d'(v_{k-1}) - L(G',R,\ell)/2$ can be made arbitrarily close to δ.

4.5 Proof of Theorem 4.1

We now outline how the above lemmas fit together to prove Theorem 4.1. For any fixed $\delta > 0$, we can pick n sufficiently large so that for good networks with at least n nodes, for the traffic rate R chosen as in Lemma 4.3, the Delta Lemmas imply that the maximum cost of travel between internal nodes is δ, and the Balance Lemmas imply that the two 'halves' of G and U are balanced to within a cost of δ.

The key idea is to fix a constant $\mu > 0$, assume that the cost along Nash flow paths of G' is $2B_2(1-\mu)$, and upper bound on the traffic rate R'. We will show that for $\mu > 0$ sufficiently small, $R' > R$. Theorem 4.1 will then follow from Proposition 2.3.

Proof of Theorem 4.1: Edges not in U will each carry at most $(B_2\mu + 2\delta)/A_0$ less flow in G' than in G, and thus these edges account for at most a discrepancy between R' and R of

$$F_1 = 2np\frac{B_2\mu + 2\delta}{A_0}.$$

Define $c = |S_1|/np$. Property (P5) implies that $q/2 \leq c \leq 1$. The edges in S_1 will each carry at most δ/A_0 units of flow in G. This accounts for at most

$$F_2 = cnp\frac{\delta}{A_0}$$

of the flow in G.

The edges in S_2 will each carry at most $(B_2+\delta)/A_2$ units of flow in G. This accounts for at most

$$F_3 = cnp\frac{B_2+\delta}{A_2}$$

of the flow in G.

Now we consider the flow through S_i and T_i in G'. Let r_1 be the minimum flow that travels along an edge of S_1 in G', and let r_2 be the minimum flow that travels along an edge of T_2 in G'. From the definition of G', the total flow along the S_1, T_2 paths must be at least $cnp\max\{r_1, r_2\}$. Furthermore, from the Delta Lemma (Lemma 4.5) we have that $r_1A_1 + B_3 + r_2A_3 + B_1 + \delta \geq 2B_2(1-\mu)$. This inequality must also hold if we replace r_1 and r_2 by $r = \max\{r_1, r_2\}$. Solving for r, we get $r \geq (2B_2(1-\mu) - B_1 - B_3 - \delta)/(A_1+A_3)$. Thus the total flow along both the S_1, T_2 paths, and the S_2, T_1 paths in G' must be at least

$$F_4 = 2cnpr \geq 2cnp\frac{2B_2(1-\mu) - B_1 - B_3 - \delta}{A_1 + A_3}.$$

Summing up the above quantities we obtain

$$\begin{aligned} R' - R &\geq F_4 - F_1 - F_2 - F_3 \\ &\geq cnp\left(\frac{2(2B_2 - B_1 - B_3)}{A_1 + A_3} - \frac{B_2}{A_2}\right) \\ &\quad -\mu cnp\left(\frac{4cB_2}{A_1 + A_3} + \frac{2B_2}{A_0}\right) \\ &\quad -\delta cnp\left(\frac{2c}{A_1 + A_3} + \frac{4+c}{A_0} + \frac{c}{A_2}\right). \end{aligned} \tag{3}$$

The parameters A_i, B_i were chosen in Subsection 4.1 so that the first term in inequality (3) is strictly positive. Thus for δ small enough (and n large enough), there is a choice of $\mu > 0$ such that $R' > R$. Combining this with Proposition 2.3 implies Theorem 4.1. $\square$

5. THE $1/X$ MODEL

We next prove an analogue of Theorem 4.1 for the $1/x$ model (see Subsection 2.2). By a random network from $\mathcal{G}(n, p, q)$, we mean a random graph G from the distribution $\mathcal{G}(n, p)$ for which each edge of G is independently given the latency function $\ell(x) = x$ with probability q and the latency function $\ell(x) = 1$ with probability $1 - q$.

Theorem 5.1 *There is a traffic rate $R = R(n, p, q)$ such that, with high probability as $n \to \infty$, the Braess ratio of a random n-node network from $\mathcal{G}(n, p, q)$ with traffic rate R is at least*

$$\frac{4 - 3pq}{3 - 2pq}.$$

Remark 5.2 Note for small values of p and q, the Braess ratio in Theorem 5.1 approaches 4/3, the worst-case bound for Braess paradox given linear latency functions [28, 30].

6. EXTENSIONS

We conclude by discussing how to relax the requirement that the $\mathcal{B}$ distribution is dense near zero, and also consider the possibility of extending the main theorem to the case where $p = o(1/\sqrt{n})$.

The condition that the $\mathcal{B}$ distribution is dense near 0 can be weakened. From inequality (3) it follows that there is some constant $C = C(\mathcal{A}, \mathcal{B})$ such that if δ is bounded above by C, then the Braess ratio of random network will be bounded away from 1 with high probability as $n \to \infty$. Provided the distribution $\mathcal{B}$ is sufficiently dense near a constant $B_0 < C/2$, we can rework our main proof to show that Theorem 4.1 still holds.

We believe, but have not verified, that our results and proof techniques can be extended to the $\mathcal{G}(n, p)$ random graph model with smaller values of p. The key technical challenge is to extend properties (P1) and (P2) of good networks—which state that various pairs of vertices are highly connected using short paths—to sparser random graphs. As the network becomes increasingly sparse, longer paths must be used to achieve the desired degree of connectivity, which in turn leads to several technical complications.

Finally, we suspect that further extensions to other random graph and latency function models are possible, but leave this for future work.

7. REFERENCES

[1] E. Altman, R. El Azouzi, and O. Pourtallier. Avoiding paradoxes in routing games. In *Proceedings of the 17th International Teletraffic Conference*, 2001.

[2] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[3] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989. Second Edition, Athena Scientific, 1997.

[4] B. Bollobás. *Random Graphs*. Academic Press, 1985.

[5] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. English translation in [6].

[6] D. Braess. On a paradox of traffic planning. *Transportation Science*, 39(4):446–450, 2005.

[7] S. C. Dafermos and A. Nagurney. On some traffic equilibrium theory paradoxes. *Transportation Research, Series B*, 18(2):101–110, 1984.

[8] S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, 73(2):91–118, 1969.

[9] R. El Azouzi, E. Altman, and O. Pourtallier. Properties of equilibria in competitive routing with several user types. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 4, pages 3646–3651, 2002.

[10] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.

[11] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1):21–41, 2001. Preliminary version in *STOC '00*.

[12] C. Fisk and S. Pallottino. Empirical evidence for equilibrium paradoxes with implications for optimal planning strategies. *Transportation Research, Part A*, 15(3):245–248, 1981.

[13] M. Frank. The Braess Paradox. *Mathematical Programming*, 20(3):283–302, 1981.

[14] M. Frank. Cost-deceptive links on ladder networks. *Methods of Operations Research*, 45:75–86, 1983.

[15] E. J. Friedman. Genericity and congestion control in selfish routing. In *Proceedings of the 43rd Annual IEEE Conference on Decision and Control (CDC)*, pages 4667–4672, 2004.

[16] M. A. Hall. Properties of the equilibrium state in transportation networks. *Transportation Science*, 12(3):208–216, 1978.

[17] H. Kameda. How harmful the paradox can be in the Braess/Cohen-Kelly-Jeffries networks. In *Proceedings of the 21st INFOCOM Conference*, volume 1, pages 437–445, 2002.

[18] G. Kolata. What if they closed 42nd Street and nobody noticed? *New York Times*, page 38, December 25, 1990.

[19] Y. A. Korilis, A. A. Lazar, and A. Orda. Capacity allocation under noncooperative routing. *IEEE Transactions on Automatic Control*, 42(3):309–325, 1997.

[20] Y. A. Korilis, A. A. Lazar, and A. Orda. Avoiding the Braess paradox in noncooperative networks. *Journal of Applied Probability*, 36(1):211–222, 1999.

[21] L. Libman and A. Orda. The designer's perspective to atomic noncooperative networks. *IEEE/ACM Transactions on Networking*, 7(6):875–884, 1999.

[22] H. Lin, T. Roughgarden, and É. Tardos. A stronger bound on Braess's Paradox. In *Proceedings of the 15th Annual Symposium on Discrete Algorithms*, pages 333–334, 2004.

[23] H. Lin, T. Roughgarden, É. Tardos, and A. Walkover. Braess's Paradox, Fibonacci numbers, and exponential inapproximability. In *Proceedings of the 32nd Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3580 of *Lecture Notes in Computer Science*, pages 497–512, 2005.

[24] J. D. Murchland. Braess's paradox of traffic flow. *Transportation Research*, 4(4):391–394, 1970.

[25] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, 2001. Preliminary version in *STOC '99*.

[26] E. I. Pas and S. L. Principio. Braess' paradox: Some new insights. *Transportation Research, Series B*, 31(3):265–276, 1997.

[27] C. M. Penchina. Braess paradox: Maximum penalty in a minimal critical network. *Transportation Research, Series A*, 31(5):379–388, 1997.

[28] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42d Annual Symposium on Foundations of Computer Science (FOCS)*, pages 472–481, 2001. Full version to appear in *Journal of Computer and Systems Sciences*.

[29] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[30] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002. Preliminary version in *FOCS '00*.

[31] R. Steinberg and W. I. Zangwill. The prevalence of Braess' Paradox. *Transportation Science*, 17(3):301–318, 1983.

[32] A. Taguchi. Braess' paradox in a two-terminal transportation network. *Journal of the Operations Research Society of Japan*, 25(4):376–388, 1982.

[33] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.