

FIRST-PASSAGE PERCOLATION ON A WIDTH-2 STRIP, AND THE PATH COST IN A VCG AUCTION

ABRAHAM FLAXMAN, DAVID GAMARNIK, AND GREGORY B. SORKIN

ABSTRACT. This paper studies the time constant for first-passage percolation and the Vickery-Clarke-Groves (VCG) payment for the shortest path on a width-2 strip with random edge costs. These statistics attempt to describe two seemingly unrelated phenomena, arising in physics and economics respectively: the first-passage percolation time predicts how long it takes for a fluid to spread through a random medium, while the VCG payment for the shortest path is the cost maximizing social welfare among selfish agents. However, our analyses of the two are quite similar, and require solving (slightly different) recursive distributional equations.

For first-passage percolation where the edge costs are independent Bernoulli random variables we find the time constant exactly. It turns out to be a rational function of the Bernoulli parameter; for example, for $\text{Be}(1/2)$ the time constant is $9/28$.

For first-passage percolation where the edge costs are uniform random variables we present a reasonably efficient means for obtaining upper and lower bounds; our calculations show the time constant to be between 0.42 and 0.43. Using Harris chains we also show that the incremental cost to advance through the medium has a unique stationary distribution, and we compute stochastic lower and upper bounds. We rely on no special properties of the uniform distribution; the same methods could be applied to any well-behaved, bounded distribution.

For the VCG payment, we restrict our attention to the case where the edge costs are independent Bernoulli random variables. Here we find that the VCG payment for a strip of length n tends to n times a rational function of the Bernoulli parameter, which we determine explicitly. For $\text{Be}(1/2)$, the VCG payment over n tends to $1829/1568 \approx 1.166$, which is about 3.63 times the length of the shortest path.

1. INTRODUCTION

The general topic of this paper is the random structure produced when a fixed graph is assigned edge costs independently at random. We will focus on a particular fixed graph, the n -long width-2 strip (defined below), and study some aspects of a minimum-cost path. In particular, we will consider the time constant for first-passage percolation, and the Vickery-Clarke-Groves (VCG) payment. These statistics attempt to describe two seemingly unrelated phenomena arising in physics and economics, respectively. However, our analyses of the two are quite similar.

First-passage percolation: First-passage percolation is a model of the time it takes a fluid to spread through a random medium [BH57, HW65, Kes87]. Mathematically, it is described by the shortest edge-weighted paths from an origin to every other point in a graph. For our purposes, the “time constant” is the limiting ratio of this length to the unweighted shortest path length n , as n tends to infinity. Previous research has derived upper and lower bounds for the time constant of first-passage percolation on the grid [SW78, Jan81, AP02] and on the random graph $G_{n,p}$ [HHM01]. For the width-2 strip, we provide a method of calculating the rate constant for any discrete edge-length distribution; the method can also be used to provide arbitrarily good bounds for any continuous distribution, as we illustrate for the uniform distribution on $[0, 1]$. Our method is similar in spirit to

the Objective Method (or Local Weak Convergence) [Ald01, AS04, GNS, Ald05], in the sense that rather than looking at a collection of random variables whose cardinality diverges to infinity as the size of the underlying graph diverges to infinity, we look at a relative quantity defined on a localized part of the graph. Previously this has been achieved by looking at the relative value of an edge included in the potential largest matching or similar structure. The corresponding quantity is called a *bonus*, and a recursive distributional equation defining the distribution of the bonus is derived. (See also the survey of recursive distributional equations [AB05].) The model in the present paper is considerably simpler due to the structure of the width-2 strip, which makes the underlying recursive distributional equation simply a Markov chain.

Because it is a Markov chain, the analysis for discrete edge-length distributions is straightforward; for a Bernoulli edge-length distribution $\text{Be}(p)$ the incremental cost $\gamma(n)$ to go from stage $i - 1$ to i has a unique stationary distribution with a simple, closed-form expression. When the edge-length distribution is continuous (uniform, for example), replacing it with a rounded-down (respectively, rounded-up) discretized equivalent gives a lower (resp., upper) bound on the expectation of the incremental cost (the time constant), but gives no information about its distribution. A subtly different approach gives stochastic lower and upper bound bounds on the incremental cost, and, separately, an analysis via *Harris chains* shows it to have a unique stationary distribution. The Harris-chain approach is well known in probability theory, but is worthy of greater attention in tangential fields.

VCG Payment: The Vickery-Clarke-Groves (VCG) mechanism [Vic61, Cla71, Gro73] applies to a setting in economics where each edge of a graph is controlled by a different selfish agent, and each agent has some private value describing the cost of using her edge. Anyone interested in buying a path in such a network is faced with the problem that an agent will lie about her edge cost if such a lie will yield her a higher payment. The VCG mechanism provides a solution to this problem in which payments to agents are carefully selected to produce the cheapest path (maximizing social welfare) in such a way that each agent finds it in her best interest to reveal her true edge cost. The VCG mechanism was first applied to the shortest-path problem explicitly in [NR99].

Unfortunately, the VCG mechanism may result in the auctioneer paying much more than the cost of the shortest path, and the overpayment can be large. Even in the case where the second-best path has cost close to that of the best path, the VCG overpayment can be large. See [AT02] for a detailed study of the worst-case behavior of this overpayment. Additional investigation of shortest paths in this setting appear in [MPS03, ESS04, CR04, Elk05].

It is possible that the worst-case bounds on the cost of the VCG mechanism are overly pessimistic. To investigate this, we compare the cost of the VCG mechanism with the shortest-path cost in the average-case setting (for the width-2 strip with random edge costs). Other average-case studies for completely different graphs appear in [MPS03, CR04], and real-world measurements appear in [FPSS02].

Generalizations: We rely on no special properties of the uniform distribution; the methods we use to analyze this edge-length distribution could equally well be applied to any well-behaved, bounded distribution.

For the $2 \times n$ strip, we show that it is not important whether edges parallel to the long direction must be traversed left-to-right or whether they can be traversed in either direction. Even for the $3 \times n$ strip, however, the distinction is important. For any fixed m , our methods apply to the

$m \times n$ strip in the left-to-right model (with some more complicated recursive equations replacing (1) and (2)), but not to the undirected model.

2. THE MODEL

Departing slightly from the usual convention, let $[n]$ denote $\{0, 1, \dots, n-1\}$. Define the *infinite width-2 strip* to be the infinite graph whose vertex set is $[2] \times \mathbb{Z}$, and whose edges join vertices at Hamming distance 1, i.e., edges (j, i) and (j', i') where $(|j-j'|, |i-i'|)$ is either $(0, 1)$ or $(1, 0)$. The *half-infinite strip* is the subgraph induced by $[2] \times \mathbb{Z}^{0,+}$, and an *n -long strip* is the (finite) subgraph induced by $[2] \times [n+1]$.

If each edge e has a weight $w(e) \in \mathbb{R}^{0,+}$, for each vertex v let $P(v)$ be the “shortest” (minimum-weight) path from $(0, 0)$ to v , and let $\ell(v)$ be the weight of this path. We consider two models: the “general-path” (GP) model where $P_{\text{GP}}(v)$ may be any path from $(0, 0)$ to v , and the “left-right” (LR) model where $P_{\text{LR}}(v)$ is restricted to be a left-to-right path. That is, $P_{\text{LR}}(v)$ is the shortest path to v which does not traverse any edge from right to left, or, still more precisely, which contains no successive pair of vertices $(j, i), (j, i-1)$.

Suppose that the edge weights are drawn independently from some given distribution, such as $\text{Be}(p)$ (the Bernoulli distribution with parameter p : $X = 1$ with probability p ; $X = 0$ w.p. $1-p$) or $U[0, 1]$ (the uniform distribution over the interval $[0, 1]$). Our first-passage percolation problem is simply to determine, for each of three types of strips, for a given distribution, and under the general-path or left-right model, the existence and value of the limiting time constant or “rate” of percolation

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}\ell(0, n)}{n}.$$

We will also show that $\ell(0, n)/n$ almost surely converges to this value, and that both statements hold also for $\ell(1, n)$, with the same rate. Note that for all our purposes it suffices to determine path lengths up to an additive constant.


For notational convenience, for any $a \leq b$, we define $\text{trunc}(x; a, b) := \max\{\min\{x, b\}, a\}$. That is, $\text{trunc}(x; a, b)$ is the “truncation” of x to the interval $[a, b]$: x if $a \leq x \leq b$; a if $x < a$; and b if $x > b$.

3. SHORTEST PATHS

The following lemma shows that, up to an additive error of at most 2, distances to $(0, n)$ or to $(1, n)$, under any of the three graph models and the two distance models, are all equivalent.

Lemma 1. *Let G denote the infinite width-2 strip with an arbitrary, fixed set of edge weights in the range $[0, 1]$ (respectively, random i.i.d. non-negative weights with expectation ≤ 1). Let H be the half-infinite restriction of G , and, for any $n \geq 0$, let K be the n -long restriction. Then, for any $j \in [2]$ and $i \in [n+1]$, the distances (respectively, expected distances) $\ell_{\text{LR}}(j, i)$ and $\ell_{\text{GP}}(j, i)$, measured in the three graphs G , H , and K , span a range of at most 2.*

Proof. We will argue only the case of fixed edge weights; the random case proceeds identically. The cheapest GP path in G from $(0, 0)$ to whichever of $(0, i)$ and $(1, i)$ is cheaper is at most as expensive as any of the paths under consideration, because this path is the least constrained; denote this path $P_{\text{GP}}^G(i)$. Fixing $j = 0$ (the $j = 1$ case is treated identically), the most constrained problem version is to find the cheapest LR path in K from $(0, 0)$ to $(0, i)$; the resulting path $P_{\text{LR}}^K(0, i)$ is

FIGURE 1. Moving from Δ_i to Δ_{i+1} .

the most expensive one under consideration. By the nature of the width-2 strip, the restriction of $P_{\text{GP}}^G(i)$ to K , unioned with the edges $\{(0,0), (1,0)\}$ and $\{(1,i), (0,i)\}$, is or includes a LR path in K from $(0,0)$ to $(0,i)$. Thus $\ell_{\text{GP}}^G(i) \leq \ell_{\text{LR}}^K(0,i) \leq \ell_{\text{GP}}^G(i) + 2$, and all the other lengths must also lie in this range. \square

Because of this, we will henceforth consider only LR paths, on the half-infinite strip H , to points $(0,n)$ and $(1,n)$. For convenience, we will write $\ell_{\text{LR}}^H(1,i)$ simply as $\ell(1,i)$ and $\ell_{\text{LR}}^H(0,i)$ as $\ell(0,i)$ or just $\ell(i)$. Define

$$\Delta(i) = \ell(1,i) - \ell(0,i).$$

With reference to Figure 1, for any $i > 0$, let X_i be the cost of the edge $\{(0,i-1), (0,i)\}$ and Y_i the cost of $\{(1,i-1), (1,i)\}$, and for any $i \geq 0$ let Z_i be the cost of $\{(0,i), (1,i)\}$.

Observe that for $i > 0$,

$$\begin{aligned} (1) \quad \gamma(i) &:= \ell(i) - \ell(i-1) = \min\{X_i, \Delta(i-1) + Y_i + Z_i\} \\ (2) \quad \Delta(i) &= \text{trunc}(\Delta(i-1) + Y_i - X_i; -Z_i, Z_i). \end{aligned}$$

Since $\Delta(i-1)$ depends only on values of X , Y , and Z with indices $i-1$ and smaller, the four random variables $\Delta(i-1)$, X_i , Y_i , and Z_i are mutually independent.

4. THE BERNOULLI CASE

Suppose that all the random variables X_i , Y_i , and Z_i are i.i.d. with distribution $\text{Be}(p)$, i.e., each is 1 with probability p and 0 w.p. $1-p$.

A “stationary distribution” for equation (2) is a distribution for $\Delta(i-1)$ giving rise to $\Delta(i)$ with the same distribution (though not necessarily independent).

Lemma 2. *When the edge weights are i.i.d. with distribution $\text{Be}(p)$, $0 \leq p < 1$, $\Delta(i)$ is a Markov chain on $\{-1, 0, 1\}$ with a unique stationary distribution, namely $\Delta = 1$ w.p. q ; $\Delta = -1$ w.p. q ; and $\Delta = 0$ w.p. $1 - 2q$, where $q = \frac{p^2}{1+3p^2}$.*

Proof. All values in question are integral, and each $\Delta(i) \leq 1$, since $(1, i+1)$ may at worst be reached via $(0, i+1)$ at an additional cost of at most 1. Symmetrically, each $\Delta(i) \geq -1$. By the independence of $\Delta(i-1)$ from (X_i, Y_i, Z_i) , $\Delta(i)$ is a Markov chain on the state space $\{-1, 0, 1\}$.

From (2), if $Z_i = 0$ (which occurs w.p. p) then $\Delta(i) = 0$. Otherwise we have the following table of possibilities, their probabilities (including the probability p that $Z_i = 1$), and the corresponding values of $\Delta(i)$:

$\Delta(i-1)$	X_i	Y_i	\mathbb{P}	$\Delta(i)$
1	0	0	$pq \cdot (1-p)^2$	1
1	0	1	$pq \cdot p(1-p)$	1
1	1	0	$pq \cdot p(1-p)$	0
1	1	1	$pq \cdot p^2$	1
0	0	0	$p(1-2q) \cdot (1-p)^2$	0
0	0	1	$p(1-2q) \cdot p(1-p)$	1
0	1	0	$p(1-2q) \cdot p(1-p)$	-1
0	1	1	$p(1-2q) \cdot p^2$	0
1	0	0	$pq \cdot (1-p)^2$	-1
-1	0	1	$pq \cdot p(1-p)$	0
-1	1	0	$pq \cdot p(1-p)$	-1
-1	1	1	$pq \cdot p^2$	-1

If $\Delta(i-1) = 1$ and $\Delta(i) = 1$ are both to have probability q , we must have

$$q = pq \cdot (1-p)^2 + pq \cdot p(1-p) + pq \cdot p^2 + p(1-2q) \cdot p(1-p),$$

whose solution is $q = p^2/(1+3p^2)$. Thus if Δ is to be stationary, we must have, for this value of q , $\Delta = 1$ w.p. q ; by symmetry $\Delta = -1$ w.p. q ; and thus $\Delta = 0$ w.p. $1-2q$.

The Markov chain's transition matrix, which corresponds to the table above (plus the 12 omitted cases when $Z_i = 0$), is easily seen to be ergodic and aperiodic as long as $0 < p < 1$, and thus has a unique stationary distribution. When $p = 0$, $\Delta_i = 0$, deterministically, for all $i \geq 0$, which still happens to fit the same form. (When $p = 1$, $\Delta_i = 1$ deterministically: the sole exception.) \square

Lemma 3. *When the edge weights are i.i.d. random variables with distribution $\text{Be}(p)$, $0 < p < 1$, $\gamma(i) = \ell(i) - \ell(i-1)$ is a Markov chain on $\{-1, 0, 1\}$ with a unique stationary distribution: it is -1 w.p. $p^2(1-p)^2/(3p^2+1)$; 1 w.p. $2p^2(1+p^2)/(3p^2+1)$; and 0 with the remaining probability, giving $\mathbb{E}(\gamma(i)) = p^2(1+p)^2/(3p^2+1)$.*

Proof. That $\gamma(i)$ is a Markov chain, and is ergodic and aperiodic, follows as in the proof of the preceding lemma. Since $\gamma(i)$ depends on four independent random values all of whose distributions are known, calculating it is straightforward. Instead of presenting a table as above we divide it into a few cases. It is -1 iff $\Delta(i) = -1$, $Y_i = 0$, and $Z_i = 0$ (the value of X_i is irrelevant), which occurs w.p. $q(1-p)^2$. It is 1 iff $X_i = 1$ and $\Delta(i) + Y_i + Z_i \geq 1$, the latter of which is satisfied if $\Delta(i) = -1$ and $Y_i = Z_i = 1$, if $\Delta(i) = 0$ and (Y_i, Z_i) is anything but $(0, 0)$, or if $\Delta(i) = 1$, giving total probability $p[qp^2 + (1-2q)(1 - (1-p)^2) + q]$. The rest of the calculation is trivial. \square

Theorem 4. *When the edge weights are i.i.d. $\text{Be}(p)$ random variables, for any p with $0 < p < 1$, $\lim_{n \rightarrow \infty} \frac{\mathbb{E}\ell(n)}{n} = p^2(1+p)^2/(3p^2+1)$, and almost surely, $\lim_{n \rightarrow \infty} \frac{\ell(n)}{n}$ exists and has the same value.*

Proof. We have established that $\gamma(i)$ is an ergodic Markov chain with the unique stationary distribution described in Lemma 3. The ergodicity implies that almost surely

$$\lim_{n \rightarrow \infty} \frac{\ell(n)}{n} = \lim_{n \rightarrow \infty} \sum_{1 \leq i \leq n} \frac{\ell(i) - \ell(i-1)}{n} = \lim_{n \rightarrow \infty} \sum_{1 \leq i \leq n} \frac{\gamma(i)}{n} = \lim_{n \rightarrow \infty} \mathbb{E}(\gamma(i)) = p^2(1+p)^2/(3p^2+1).$$

Since the values $\gamma(i)$ are bounded almost surely (in fact surely, by unity, in absolute value), the almost sure convergence implies the convergence in expectation. \square

5. UNIFORM CASE: EXPECTATION

What if $X_i, Y_i,$ and Z_i have uniform distribution, $U[0, 1]$? The development in the previous sections applies here as well: $\gamma(i)$ and $\Delta(i)$ are again Markov chains, but with continuous state space. Rather than working with the continuous state space we introduce a natural and simple discretization. Fix a (large) positive integer k and let \bar{U}_k (respectively \underline{U}_k) be the uniform distribution on the set $\{1/k, \dots, (k-1)/k, 1\}$ (respectively $\{0, 1/k, \dots, (k-1)/k\}$). Lemma 5 establishes the simple observation that the length $\ell(n)$ is bounded from above and below by the lengths corresponding to these discrete uniform distributions. This enables us to circumvent infinite state spaces (the discretized Markov chain is finite) and to compute upper and lower bounds on the limiting value of $\ell(n)/n$.

Lemma 5. *Let $\ell(i)$, $\bar{\ell}_k(i)$, and $\underline{\ell}_k(i)$ be the length of the shortest path when the edge weight distribution is $U[0, 1]$, \bar{U}_k , and \underline{U}_k respectively. Then the limits $\lim_{n \rightarrow \infty} \frac{\mathbb{E}\underline{\ell}_k(n)}{n}$ and $\lim_{n \rightarrow \infty} \frac{\mathbb{E}\bar{\ell}_k(n)}{n}$ exist and*

$$(3) \quad \begin{aligned} \lim_{n \rightarrow \infty} \frac{\mathbb{E}\underline{\ell}_k(n)}{n} &\leq \liminf_{n \rightarrow \infty} \frac{\mathbb{E}\ell(n)}{n} \\ &\leq \limsup_{n \rightarrow \infty} \frac{\mathbb{E}\ell(n)}{n} \leq \lim_{n \rightarrow \infty} \frac{\mathbb{E}\bar{\ell}_k(n)}{n} \leq \lim_{n \rightarrow \infty} \frac{\mathbb{E}\underline{\ell}_k(n)}{n} + 2/k. \end{aligned}$$

Proof. Suppose the weight distribution is \bar{U}_k . Since the values of $\Delta(0) = \ell(1, 0) - \ell(0, 0) = Z_0$ have support $\{1/k, \dots, k/k\}$, it is straightforward to check from the recursion (1) and (2) that the values of $\Delta(i)$ and $\gamma(i) = \ell(i) - \ell(i-1)$ have support $\{-1, -(k-1)/k, \dots, (k-1)/k, 1\}$. The same recursions imply that $\gamma(i)$ is a Markov chain on this (finite) state space. We claim that this Markov chain is ergodic and has a unique stationary distribution. For this it suffices to show that for some particular state, there is positive probability of making a transition from any state to the particular state. Indeed, for every value of Δ_{i-1} let $X_i = 1/k$ and $Y_i = Z_i = 1$. Then $\Delta_{i-1} + Y_i + Z_i \geq 1 > X_i$, and (1) gives $\Delta_i = X_i = 1/k$. Thus the state $1/k$ is reachable with positive probability from any other state and the ergodicity is established. Repeating the discussion in previous section, we obtain that the limit $\lim_{n \rightarrow \infty} \bar{\ell}_k(n)/n$ exists in expectation and almost surely.

A similar analysis applies for the distribution \underline{U}_k . For any value $x \in [0, 1]$, define \bar{x} to be the value of x rounded up to the nearest multiple of $1/k$ (i.e., $\bar{x} = \lceil kx \rceil / k$), and \underline{x} to be the similar rounding-down. We observe that if we replace the values $X_i, Y_i,$ and Z_i with $\bar{X}_i, \bar{Y}_i,$ and \bar{Z}_i then the value of $\ell(n)$ increases, and at the same time the distribution $U[0, 1]$ is transformed into \bar{U}_k . Similarly, if we round $X_i, Y_i,$ and Z_i down to $\underline{X}_i, \underline{Y}_i,$ and \underline{Z}_i then we decrease the value of $\ell(n)$ and transform $U[0, 1]$ into \underline{U}_k . Finally, note that $\bar{X}_i \leq X_i + 1/k$ and similar bounds hold for Y and Z . Since each term $\ell(i) - \ell(i-1)$ uses at most two of these random variables, $\bar{\ell}_k(n) \leq \underline{\ell}_k(n) + n(2/k)$. This establishes the validity of the bounds (3) and completes the proof of the lemma. \square

Corollary 6. *The limit $\lim_{n \rightarrow \infty} \frac{\mathbb{E}\ell(n)}{n}$ exists, and*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{E}\bar{\ell}_k(n)/n = \lim_{n \rightarrow \infty} \frac{\mathbb{E}\ell(n)}{n} = \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{E}\underline{\ell}_k(n)/n.$$

We have used the bounding distributions \underline{U}_k and \bar{U}_k to compute the expected value numerically. We exploit the simple observation that $\lim_{n \rightarrow \infty} \mathbb{E}[\bar{\ell}_k(n)/n] = \lim_{i \rightarrow \infty} \mathbb{E}[\bar{\gamma}_k(i)]$ (where of course we define $\bar{\gamma}_k(i) = \bar{\ell}_k(i) - \bar{\ell}_k(i-1)$). We also remark that, computationally, it suffices first to compute




FIGURE 2. Density functions for the stationary distributions of $\underline{\gamma}_{190}(i)$ and $\overline{\gamma}_{190}(i)$, giving bounds of $0.420 < \mathbb{E}[\gamma(i)] < 0.431$. The discrete density functions are plotted as continuous lines for visual clarity.

the stationary distribution for $\overline{\Delta}_k$ by finding the fixed point for (2), from which the stationary distribution for $\overline{\ell}_k$ follows by a single application of (1).

Taking $k = 190$, we find the percolation rate for the upper discretized uniform distribution, $\lim_{i \rightarrow \infty} \mathbb{E}[\overline{\gamma}_k(i)]$, to be just below 0.4306. A corresponding value of just above 0.4200 for $\underline{\ell}_{190}$ means that the true percolation rate lies between 0.420 and 0.431. The same calculation for values $k > 190$ yields arbitrarily good bounds on the rate.

Note that the direct path from $(0, 0)$ to $(0, n)$ has expected length exactly $0.5n$. Thus the average “saving” rendered by the availability of the alternative edges Y_i, Z_i is between $0.069n$ and $0.080n$.

Also with $k = 190$, the density of the stationary distribution (the $i \rightarrow \infty$ limit) of $\overline{\gamma}_k(i)$ is plotted in Figure 2. Unfortunately there is no significance of this for other values of k : in principle the distribution for $k = 191$ could look completely different. We address this and related issues in the next section. Meanwhile, the observation of the computed discrete density’s unimodality and its large jump at zero prompts a corresponding conjecture for the continuous version.

Conjecture 7. *Under the uniform distribution $U[0, 1]$, the density of the stationary distribution of $\gamma(i)$ is continuous everywhere but at 0, and is unimodal with mode approximately 0.23.*

We note that the method of using rounded-up and rounded-down versions of the random edge costs applies to an arbitrary bounded distribution as well as to the uniform distribution.

6. UNIFORM CASE: DISTRIBUTION

In the previous section we derived random variables $\underline{\gamma}(n)$ and $\overline{\gamma}(n)$ whose expectations bounded that of the quantity of interest, $\mathbb{E}[\underline{\gamma}(n)] \leq \mathbb{E}[\gamma(n)] \leq \mathbb{E}[\overline{\gamma}(n)]$, and whose limiting expectations could be computed. Unfortunately, the random variables tell us nothing about the distribution of

$\gamma(n)$. For example, it is not generally true that $\underline{\gamma}(n) \leq \gamma(n)$, because $\underline{\gamma}(n)$ is a difference between two rounded-down quantities, and thus may be larger than the difference between the unrounded quantities.

In this section we define different random variables (which we will also call $\underline{\gamma}(n)$ and $\overline{\gamma}(n)$ because they play a similar role), which *do* provide stochastic lower and upper bounds on $\gamma(n)$: there is a coupling (in this case the natural coupling) in which it is always true that $\underline{\gamma}(n) \leq \gamma(n) \leq \overline{\gamma}(n)$. While the motive behind the construction is quite different, the calculations are almost identical.

We also argue from ‘‘Harris chains’’ that the random variables $\gamma(n)$ have a limiting distribution. Calculating the $n \rightarrow \infty$ limiting distributions of $\underline{\gamma}_k(n)$ and $\overline{\gamma}_k(n)$ (for example with $k = 150$) provides bounds on the limiting distribution of $\gamma(n)$.

We work entirely with the Δ of recurrence (2), from which corresponding conclusions for γ follow immediately via (1). In all the following, by a ‘‘lower bound’’ we mean a stochastic lower bound. In the previous section, $\underline{\Delta}_k(n)$ (we will now drop the argument n) denoted an (exact) random variable corresponding to a Markov chain on values such as \underline{X}_k ; \underline{X}_k is a lower bound on X_k , but $\underline{\Delta}_k$ was not necessarily a lower bound on Δ . In this section, we redefine $\underline{\Delta}_k$ (or more generally $\underline{\Delta}$) so that it *is* a lower bound on Δ , but no longer corresponds to the discretized Markov chain on \underline{X}_k etc.

Letting $W = Y - X$, from (2),

$$\begin{aligned} \Delta' &= \text{trunc}(\Delta + (Y - X); -Z, Z) \\ &= \text{trunc}(\Delta + W; -Z, Z) \\ &\geq \text{trunc}(\underline{\Delta} + \underline{W}; -\underline{Z}, \underline{Z}), \end{aligned}$$

where \underline{Z} and \overline{Z} are any random variables with the property that $\underline{Z} \leq Z \leq \overline{Z}$. Choosing \underline{Z} and \overline{Z} , as before, to be rounded-down and rounded-up discretizations of Z , and for convenience defining $\epsilon = 1/k$,

$$\begin{aligned} \Delta' &\geq \text{trunc}(\underline{\Delta} + \underline{W}; -\underline{Z} - \epsilon, \underline{Z}) \\ &= \underline{W} + \text{trunc}(\underline{\Delta}; -\underline{Z} - \underline{W} - \epsilon, \underline{Z} - \underline{W}) \\ (4) \quad &=: \underline{\Delta}'. \end{aligned}$$

Similarly,

$$(5) \quad \Delta' \leq \underline{W} + \epsilon + \text{trunc}(\overline{\Delta}; -\underline{Z} - \underline{W} - \epsilon, \underline{Z} - \underline{W}) =: \overline{\Delta}'.$$

Setting $\underline{\Delta}(0) = -1$ (deterministically) and $\overline{\Delta}(1) = 1$, of course $\underline{\Delta}(0) \leq \Delta(0) \leq \overline{\Delta}(0)$. Using equation (4) to define $\underline{\Delta}(n) = \underline{\Delta}'$ from $\underline{\Delta}(n-1) = \underline{\Delta}$, and similarly using (5) to define $\overline{\Delta}(n)$ from $\overline{\Delta}(n-1)$, we have that, for all n , $\underline{\Delta}(n) \leq \Delta(n) \leq \overline{\Delta}(n)$. From (1), trivially,

$$(6) \quad \gamma(i) \geq \min\{\underline{X}_i, \underline{\Delta}(i-1) + \underline{Y} + \underline{Z}_i\}$$

$$(7) \quad \gamma(i) \leq \min\{\overline{X}_i, \overline{\Delta}(i-1) + \overline{Y} + \overline{Z}_i\}$$

Theorem 9 will show that $\gamma(n)$ itself has a unique stationary distribution. Meanwhile, for any fixed k , the Markov chains for $\underline{\Delta}(n)$ and $\overline{\Delta}(n)$ are both well-behaved finite Markov chains, with stationary distributions we will call $\underline{\Delta}$ and $\overline{\Delta}$. Substituting $\underline{\Delta}$ and $\overline{\Delta}$ into (6) and (7) defines corresponding random variables $\underline{\gamma}$ and $\overline{\gamma}$, which are then stochastic lower and upper bounds on $\gamma(n)$ for $n \geq 10$ (and for its stationary distribution). Distribution functions for $\underline{\gamma}$ and $\overline{\gamma}$ are plotted




FIGURE 3. Distribution functions for the stationary distributions of $\underline{\gamma}$ and $\overline{\gamma}$ given by $k = 150$. For any $n > 10$, the true incremental-length distribution $\gamma(n)$ lies between the two (incidentally proving that $0.4215 < \mathbb{E}\gamma < 0.4292$).

in Figure 3. By construction, the two curves never cross; the bounds are sufficiently good that they are largely visually indistinguishable. Of course, $\mathbb{E}\underline{\gamma} \leq \mathbb{E}\gamma \leq \mathbb{E}\overline{\gamma}$, and as it happens these bounds are better than the previous ones despite the use of $k = 150$ here and $k = 190$ there: $0.4215 < \mathbb{E}\gamma < 0.4292$.

Remark: Since $\underline{\Delta}$ and \underline{W} each have support of cardinality $2k+1$ and \underline{Z} has support of cardinality k , the $(2k+1) \times (2k+1)$ transition matrix for $\underline{\Delta}$ can be constructed in time $O(k^3)$. The stationary distribution of $\underline{\Delta}$ is the eigenvector of the transition matrix corresponding to eigenvalue 1, and can be computed in time $O(k^3)$ by row reduction. Instead, we approximated the eigenvector by taking the initial distribution vector $\underline{\Delta}(0)$ and repeatedly multiplying it by the transition matrix; while we made no effort to compute the convergence rate, with $k = 150$ a stationary vector $\underline{\Delta}$ was obtained (to within machine precision) in under 10 iterations. The same statements apply to $\overline{\Delta}$, \overline{W} and \overline{Z} .

Just as (3) showed that our upper and lower bounds on the expectation were within $2/k$ of one another, here we can prove that, for any k , the stationary random variables $\underline{\Delta}$ and $\overline{\Delta}$ are also arbitrarily near to one another: $d(\underline{\Delta}, \overline{\Delta}) = O(1/k)$, where we define the distance between continuous random variables X and Y as the area between their CDFs (cumulative density functions). (For any coupling of two variables X and Y , $\mathbb{E}[|X - Y|] \geq d(X, Y)$, with an optimal coupling giving equality.) Recall that $\epsilon = 1/k$.

Theorem 8. *The stationary random variables $\underline{\Delta}$ and $\overline{\Delta}$ for equations (4) and (5) satisfy $d(\underline{\Delta}, \overline{\Delta}) = O(\epsilon)$.*

Proof. Let $\underline{\Delta}$ and $\overline{\Delta}$ be the values of the coupled random variables at some time, defined by their previous values and random variables W and Z , and let $D := \overline{\Delta} - \underline{\Delta}$. Let $\underline{\Delta}'$, D' , etc. be the corresponding values one step later, and $\underline{\Delta}''$ etc. those two steps later. It may help to think of $\underline{\Delta}$, $\overline{\Delta}$, and D as the state variables of the single Markov chain representing the coupled process.

Because in the coupling the random variables $\underline{\Delta}, \overline{\Delta}$ always satisfy $\underline{\Delta} \leq \overline{\Delta}$, the distance $d(\underline{\Delta}, \overline{\Delta})$ between the distributions is by definition $\mathbb{E}(D)$.

Condition on $D = d$. First we observe that $D' \leq d + \epsilon$, and thus $D'' \leq d + 2\epsilon$; this is immediate from (4) and (5). (The two truncations have identical lower and upper ‘‘stops’’, so their difference is at most the difference of their central arguments, and the other terms in (4) and (5) contribute a difference of exactly ϵ .)

Second, if the truncation defining $\underline{\Delta}''$ takes the value of its upper stop then so does that for $\overline{\Delta}''$, in which case $D'' \leq \epsilon$. This happens with constant probability, as we now argue. We assume $\epsilon \leq 0.01$. From the right hand side of the first line of (4), if $Z' \leq 0.1$ then

$$\underline{\Delta}' \geq \underline{Z}' - \epsilon \geq (Z' - \epsilon) - \epsilon \geq -0.12.$$

If in addition $W'' \geq 0.5$, and $Z'' \leq 0.37$, then

$$\underline{\Delta}' + \underline{W}'' \geq \underline{\Delta}' + (W'' - \epsilon) \geq 0.37 \geq Z'' \geq \underline{Z}'',$$

and from (4) $\underline{\Delta}'' = \underline{Z}''$. These events occur with probability $\mathbb{P}(Z' \leq 0.1) \cdot \mathbb{P}(W'' \geq 0.5) \cdot \mathbb{P}(Z'' \leq 0.37) = 0.1 \cdot \frac{1}{8} \cdot 0.37 > 0.001$.

Assembling these two facts, conditioned on $D = d$, with probability ≥ 0.001 $D'' \leq \epsilon$ while with the remaining probability $D'' \leq d + 2\epsilon$. Thus $\mathbb{E}(D'' \mid D = d) \leq 0.001\epsilon + 0.999(d + 2\epsilon) < 2\epsilon + 0.999d$. Integrating, $\mathbb{E}(D'') \leq 2\epsilon + 0.999\mathbb{E}(D)$. Since the Markov chain is in stationarity, $\mathbb{E}(D'') = \mathbb{E}(D)$, so $0.001\mathbb{E}(D) \leq 2\epsilon$, i.e., $\mathbb{E}(D) \leq 2000\epsilon$. \square

From (1), a lower-bounding random variable $\underline{\gamma}_i \preceq \gamma_i$ is given by $\underline{\gamma}_i = \min\{\underline{X}_i, \underline{\Delta}(i-1)\}$, and an upper-bounding one by $\overline{\gamma}_i = \min\{\overline{X}_i, \overline{\Delta}(i-1)\} \leq \min\{\underline{X}_i + \epsilon, \overline{\Delta}(i-1)\}$. In the coupling, the random variables' values always satisfy $0 \leq \overline{\gamma}_i - \underline{\gamma}_i \leq \epsilon + (\overline{\Delta}(i-1) - \underline{\Delta}(i-1))$. Taking expectations over the stationary distributions we know to exist (these are finite-state Markov chains) ($\overline{\Delta}, \underline{\Delta}, \overline{\gamma}$, and $\underline{\gamma}$ are all discrete random variables) gives $d(\underline{\gamma}, \overline{\gamma}) = \mathbb{E}(\overline{\gamma} - \underline{\gamma}) \leq \epsilon + \mathbb{E}(\overline{\Delta} - \underline{\Delta}) = O(\epsilon)$.

Finally, we show that Δ has a well-defined stationary distribution; from (1) it is then immediate that γ does as well.

Theorem 9. *The continuous Markov chain (2) has a unique stationary distribution.*

Definition 10. *A discrete time Markov chain $\Phi(t)$ with state space Ω is defined to be a recurrent Harris chain if there exist two sets $A, B \subset \Omega$ satisfying the following properties.*

- (1) $\Phi(t) \in A$ infinitely often w.p. 1.
- (2) There exists a non-zero measure ν with support contained in B such that for every $x \in A$ and $C \subset B$, $\mathbb{P}(\Phi(t+1) \in C \mid \Phi(t) = x) \geq \nu(C)$.

(See [Dur96] Section 5.6 pages 325-326 for a Harris chain, and page 329 for recurrent Harris.) One can use this definition to embed a recurrent Harris Markov chain into a modified chain with an extra state α , such that the special state α is visited infinitely often with probability 1. Then various recurrent properties of the chain can be deduced. Specifically, it can be established (see Durrett [Dur96]) that the recurrent Harris chain possesses a unique stationary distribution. Our next goal is to show that our chain $\Delta(i)$ is indeed recurrent Harris. Let $A = [-0.1, 1]$, $B = [0, 0.4]$ and let ν be uniform probability distribution on B multiplied by 0.2.

Lemma 11. *$\Delta(i)$ is a recurrent Harris chain with the sets A, B and measure ν defined above.*

Proof. First we note that the conclusion of Theorem 9 follows immediately from this lemma, since, as remarked above, any recurrent Harris chain possesses a unique stationary distribution. Now, to show that the chain is a recurrent Harris chain, we observe that when $\Delta(i) \in A$, that is $\Delta(i) \geq -0.1$, if in addition $W_{i+1} \geq 0.5$ and $Z_{i+1} \leq 0.4$, then $\Delta(i+1) = \text{trunc}(\Delta(i) + W_{i+1}; -Z_{i+1}, Z_{i+1}) = Z_{i+1}$. Let $V_{i+1} = 1\{Z_{i+1} \leq 0.4\}$. Note that, conditioned on $V_{i+1} = 1$, Z_{i+1} is distributed uniformly on $[0, 0.4]$. Let $p = \mathbb{P}(W_{i+1} \geq 0.5, V_{i+1} = 1) = 0.2$. Then for every $C \subset B$ and $x \in A$ we have

$$\begin{aligned} \mathbb{P}(\Delta(i+1) \in C \mid \Delta(i) = x) &\geq \mathbb{P}(W_{i+1} \geq 0.5, V_{i+1} = 1) \cdot \mathbb{P}(Z_{i+1} \in C \mid W_{i+1} \geq 0.5, V_{i+1} = 1) \\ &= p\mathbb{P}(Z_{i+1} \in C \mid V_{i+1} = 1) = p\mu(C) = \nu(C), \end{aligned}$$

where μ is the uniform measure on B and we define ν by $\nu(C) = p\mu(C)$. Therefore, $\Delta(i)$ satisfies condition (2) of Definition 10. We now prove condition (1), that w.p. 1 the set A is visited infinitely often. This is a simple corollary of the fact that if $Z_i \leq 0.1$ then $\Delta(i) = \text{trunc}(\Delta(i-1) + W_i; -Z_i, Z_i) \geq -0.1$, that is $\Delta(i) \in A$. Clearly this happens infinitely often w.p. 1. \square

7. AN AUCTION MODEL

Auction models for networks have received much attention recently. Suppose that in the half-infinite width-2 strip, each edge is provided by an independent agent who incurs a cost for supplying it (or for allowing us to drive over it, transmit data over it, or whatever). What we have considered so far would be the case where the agents' costs are random variables, and we pay for a cheapest path. However, in this setup, agents have an incentive to lie: their true cost is not the cost they will sensibly tell us.

A popular way to deal with potentially dishonest agents is to assume that each agent will act independently to maximize her own utility, and to design a mechanism where this behavior will result in every agent being truthful. It is within this framework that the Vickery-Clarke-Groves (VCG) mechanism operates [Vic61, Cla71, Gro73]; in a truthful fashion, the VCG mechanism finds the outcome that maximizes social welfare. For buying an (s, t) -path, the VCG mechanism is the following: The auctioneer finds a cheapest path, and, for each edge on that path, pays the corresponding agent the difference between the cost of a cheapest path avoiding the edge and the cost of a cheapest path if the edge cost were 0. (The mechanism is truthful because by inflating her cost, an agent does not affect the amount she gets paid, until the point when she inflates the price so much that her edge is no longer in a shortest path and she gets paid nothing.) The VCG mechanism was first applied to the shortest-path problem explicitly in [NR99].

Unfortunately, the VCG mechanism may result in the auctioneer paying much more than the cost of the shortest path, and the overpayment can be large. The simplest example comes from a source and sink connected by two parallel edges, one with cost 1 and the other with cost $c > 1$. The shortest path is the edge with cost 1, and the payment made to it is $c - 0 = c$; the ratio between this VCG cost and the simple shortest-path cost of 1 is unbounded if c is much larger than 1. In fact, even in the case where the second-best path has cost close to that of the best one, the VCG overpayment can be large; see [AT02] for a detailed study of the worst-case behavior of this overpayment. An example from [AT02] consists of two disjoint (s, t) -paths, with lengths L and $L(1 + \epsilon)$, and with the cheaper path containing k edges; the total payment is $L(1 + k\epsilon)$. Additional investigation of shortest paths in the worst-case setting appear in [MPS03, ESS04, CR04, Elk05].

It is natural to wonder how the cost of the VCG mechanism compares with the shortest-path cost in the average-case setting. We will study the cost on the width-2 strip with random edge weights.




FIGURE 4. The VCG cost at step i , working in from the left and the right, and assuming both sides are in stationarity.

(For average-case studies on completely different graphs, see [MPS03, CR04], and for real-world measurements see [FPSS02].)

Theorem 12. *When the edge weights are i.i.d. $\text{Be}(p)$ random variables, with $0 < p < 1$, the VCG path cost satisfies*

$$(8) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}(\ell_{\text{VCG}}(n)) = \frac{p(2 + 5p + 4p^2 + 8p^3 + 11p^4 - 3p^6 + p^8)}{(1 + 3p^2)^2}.$$

Proof. With reference to Figure 1, we compute the contribution of the i th triple of edges (X_i, Y_i, Z_i) to the expected VCG cost. Let $\omega(n)$ be any function tending to infinity much slower than n itself, i.e., with $1 \ll \omega(n) \ll n$. Note that any edge's contribution to the VCG cost is at most 3: we can circumvent any horizontal edge with a “loop” of 3 edges around it, each edge costing at most 1, and we can bypass any vertical edge at worst by going one more step to the right and traversing the next vertical edge, for a cost of at most 2. Thus the contribution of the first and last $\omega(n)$ edges to the limit is at most $6\omega(n)/n$, which tends to 0.

Now, for any i , a shortest path between $(0, 0)$ and $(0, n)$ may be found by taking the shortest paths from $(0, 0)$ to both $(0, i)$ and $(1, i)$, and also the shortest paths from $(0, n)$ to both $(0, i + 2)$ and $(1, i + 2)$, and finding the cheapest total way of joining one of the first paths to one of the second. The first two paths depend only on variables with indices less than i , and without loss of generality (up to an additive constant) we may consider their two costs to be 0 and Δ . Likewise, the second two paths depend only on variables with indices $i + 2$ or more, and their costs may be given as 0 and Δ' . For $\omega(n) < i < n - \omega(n)$, Δ and Δ' are independent random variables drawn from a distribution asymptotically equal to the stationary distribution. Thus, with reference to Figure 4, we consider the payments we must make for the edges X_i, Y_i , and Z_i , when Δ and Δ' are i.i.d. random variables drawn from the stationary distribution, and X_i, Y_i, Z_i, U_i , and V_i are i.i.d. $\text{Be}(p)$ random variables. Since, over all i , such groups (X_i, Y_i, Z_i) cover each edge exactly once (except for the single edge Z_0), the total of the expected payments for one such group is precisely the limiting expectation called for in (8).

This is a straightforward calculation. Dropping the subscripts for convenience, let $A = X + U$ be the cost of the path using X, U ; $B = X + Z + V + \Delta'$ that of the path using X, Z, V ; $C = \Delta + Y + V + \Delta'$ that using Y, V ; and $D = \Delta + Y + Z + U$ that using Y, Z, U . If we break ties in favor of lower letters (A in favor of B in favor of C in favor of D), the payment to X is

$$C(X) = \chi(\min(A, B) \leq \min(C, D)) \cdot [\min(C, D) - (\min(A, B) - X)],$$

that is, it is 0 unless the edge X is used, and then it is the cost of the cheapest path avoiding X less the cost of the cheapest path if X were 0, which in this case is the cheapest path using X , minus X . Similarly, the payment to Y is

$$C(Y) = \chi(\min(C, D) < \min(A, B)) \cdot [\min(A, B) - (\min(C, D) - Y)].$$

The payment to Z follows similarly, with slightly more complicated tie-breaking:

$$C(Z) = \chi((B < A) \vee (B \leq \min(C, D)) \vee (D < \min(A, B, C))) \cdot [\min(A, C) - (\min(B, D) - Z)].$$

Where the stationary probabilities for Δ and Δ' are written as $\mathbb{P}_\Delta(\cdot)$, and the Bernoulli probabilities as $\text{Be}(1) = p$ and $\text{Be}(0) = 1 - p$, the expected total payments for X , Y , and Z is

$$\sum_{X,Y,Z,U,V,\Delta,\Delta'} \text{Be}(X) \text{Be}(Y) \text{Be}(Z) \text{Be}(U) \text{Be}(V) \mathbb{P}_\Delta(\Delta) \mathbb{P}_\Delta(\Delta') \cdot [C(X) + C(Y) + C(Z)],$$

the sum taken over the $2^5 3^2$ possible values of the variables. This is a small finite sum of an explicit expression, and is calculated (by Mathematica) to be the value shown in expression (8). \square

A plot of the VCG cost rate $\lim_{n \rightarrow \infty} \mathbb{E} \ell_{\text{VCG}}(n)/n$, along with the corresponding shortest-path cost rate, is given in Figure 5, and Figure 6 shows the ratio of the VCG cost to the shortest-path cost.




FIGURE 5. VCG and usual shortest-path rates.

REFERENCES

- [AB05] David J. Aldous and Antar Bandyopadhyay, *A survey of max-type recursive distributional equations*, *Annals of Applied Probability* **15** (2005), no. 2, 1047–1110.
- [Ald01] David J. Aldous, *The $\zeta(2)$ limit in the random assignment problem*, *Random Structures Algorithms* **18** (2001), no. 4, 381–418.
- [Ald05] David J. Aldous, *Cost-volume relationships for flows through a disordered network*, ArXiv Condensed Matter e-prints, Feb 2005.
- [AP02] Sven Erick Alm and Robert Parviainen, *Lower and upper bounds for the time constant of first-passage percolation*, *Comb. Probab. Comput.* **11** (2002), no. 5, 433–445.
- [AS04] David J. Aldous and J. Michael Steele, *The objective method: probabilistic combinatorial optimization and local weak convergence*, *Probability on discrete structures*, *Encyclopaedia Math. Sci.*, vol. 110, Springer, Berlin, 2004, pp. 1–72.
- [AT02] Aaron Archer and Éva Tardos, *Frugal path mechanisms*, *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms* (San Francisco, California), January 06-08 2002, pp. 991–999.




FIGURE 6. Ratio of VCG cost to shortest-path cost.

- [BH57] S. R. Broadbent and J. M. Hammersley, *Percolation processes i. crystals and mazes*, Proceedings of the Cambridge Philosophical Society **53** (1957), 629–641.
- [Cla71] E. Clarke, *Multipart pricing of public goods*, Public Choice **8** (1971), 17–33.
- [CR04] Artur Czumaj and Amir Ronen, *On the expected payment of mechanisms for task allocation (extended abstract)*, Proceedings of the fifth ACM Conference on Electronic Commerce, 2004, pp. 252–253.
- [Dur96] Richard Durrett, *Probability: theory and examples*, second ed., Duxbury Press, Belmont, CA, 1996.
- [Elk05] Edith Elkind, *True costs of cheap labor are hard to measure: edge deletion and VCG payments in graphs*, Proceedings of the sixth ACM Conference on Electronic Commerce, 2005, pp. 108–117.
- [ESS04] Edith Elkind, Amit Sahai, and Ken Steiglitz, *Frugality in path auctions*, Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms, 2004, pp. 694–702.
- [FPSS02] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker, *A BGP-based mechanism for lowest-cost routing*, PODC '02: Proceedings of the twenty-first annual symposium on Principles of distributed computing (New York, NY, USA), ACM Press, 2002, pp. 173–182.
- [GNS] D. Gamarnik, T. Nowicki, and G. Swirszcz, *Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method*.
- [Gro73] T. Groves, *Incentives in teams*, Econometrica **41** (1973), no. 4, 617–631.
- [HHM01] Remco Van Der Hofstad, Gerard Hooghiemstra, and Piet Van Mieghem, *First-passage percolation on the random graph*, Probab. Eng. Inf. Sci. **15** (2001), no. 2, 225–237.
- [HW65] J. M. Hammersley and J. D. Welsh, *First-passage percolation, subadditive processes, stochastic networks, and generalized renewal theory*, Bernoulli-Bayes-Laplace Anniversary Volume, Springer-Verlag, 1965, pp. 61–110.
- [Jan81] S. Janson, *An upper bound for the velocity of first-passage percolation*, Journal of Applied Probability **18** (1981), 256–262.
- [Kes87] H. Kesten, *Percolation theory and first-passage percolation*, Ann. Probab. (1987), no. 15, 1231–1271.
- [MPS03] Milena Mihail, Christos Papadimitriou, and Amin Saberi, *On certain connectivity properties of the internet topology*, FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (Washington, DC, USA), IEEE Computer Society, 2003, p. 28.
- [NR99] Noam Nisan and Amir Ronen, *Algorithmic mechanism design (extended abstract)*, Proceedings of the thirty-first annual ACM symposium on Theory of computing (Atlanta, Georgia, United States), 1999, pp. 129–140.

- [SW78] R.T. Smythe and John C. Wierman, *First-passage percolation on the square lattice*, Lecture Notes in Mathematics, vol. 671, Springer-Verlag, 1978.
- [Vic61] W. Vickery, *Counterspeculation, auctions, and competitive sealed tenders*, Journal of Finance **16** (1961), 8–37.

(Abraham Flaxman) DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA

E-mail address: `abie@cmu.edu`

(David Gamarnik) SLOANE SCHOOL OF MANAGEMENT, MIT,, CAMBRIDGE, MA

E-mail address: `gamarnik@mit.edu`

(Gregory B. Sorkin) DEPARTMENT OF MATHEMATICAL SCIENCES, IBM T.J. WATSON RESEARCH CENTER, YORK-TOWN HEIGHTS, NY

E-mail address: `sorkin@watson.ibm.com`