First-Passage Percolation on a Width-2 Strip and the Path Cost in a VCG Auction

Abraham Flaxman[1], David Gamarnik[2], and Gregory B. Sorkin[3]

[1] Department of Mathematical Sciences, Carnegie Mellon University
Pittsburgh, PA
abie@cmu.edu

[2] MIT Sloan School of Management
Cambridge, MA
gamarnik@mit.edu

[3] Department of Mathematical Sciences, IBM T.J. Watson Research Center
Yorktown Heights, NY
sorkin@watson.ibm.com

Abstract. We study both 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 of maximizing social welfare among selfish agents. However, our analyses of the two are quite similar, and require solving (slightly different) recursive distributional equations. Using Harris chains, we can characterize distributions, not just expectations.

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 easier case of the width-2 strip, we provide a method of exactly calculating the time constant for any discrete edge-length distribution; the method can also be used to provide arbitrarily

P. Spirakis et al. (Eds.): WINE 2006, LNCS 4286, pp. 99–111, 2006.

good bounds for any well-behaved continuous distribution, as we illustrate for the uniform distribution on $[0, 1]$. Our method is similar in spirit to to the Objective Method (or Local Weak Convergence) [Ald01, GNSar, AS04], and is also based on constructing a certain recursive distributional equation. 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 $\mathrm{Be}(p)$ the incremental cost $\gamma(n)$ to go from stage $n-1$ to n has a unique stationary distribution with a simple, closed-form expression, and its expectation is the time constant in question. 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 time constant, but no information about the incremental cost $\gamma(n)$. 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 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 [Vic61, Cla71, Gro73]. 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 structured to yield a cheapest path (maximizing social welfare) and so 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 pay more than the cost of the shortest path, and the overpayment can be large. The VCG overpayment can be large even in the case where the second-best path has cost close to that of the best path. See [AT02] for a detailed study of the worst-case behavior of the overpayment. Additional investigations 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, KN05], 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 $m \geq 3$, our methods apply to the $m \times n$ strip in the left-to-right model (with 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 $\{0, 1, \ldots, n-1\}$. 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 join (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 $[2] \times [n+1]$.

Let each edge e have a non-negative real weight $w(e)$. 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 $P_{\mathrm{GP}}(v)$ may be any path from $(0, 0)$ to v, and the "left-right" (LR) model where $P_{\mathrm{LR}}(v)$ is restricted to be a left-to-right path. That is, $P_{\mathrm{LR}}(v)$ 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 $\mathrm{Be}(p)$ (the Bernoulli distribution with parameter p, where $X = 1$ with probability p and $X = 0$ 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,

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

We will also show that $\ell(0, n)/n$ almost surely converges to this value, and that the same statements hold 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 convenience, for $a \leq b$, define $\mathrm{trunc}(x;\ a,\ b) := \max\{\min\{x, b\}, a\}$. Thus, $\mathrm{trunc}(x;\ a,\ b)$ 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 $[0, 1]$ (resp., 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 (resp., expected distances) $\ell_{\mathrm{LR}}(j, i)$ and $\ell_{\mathrm{GP}}(j, i)$, 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

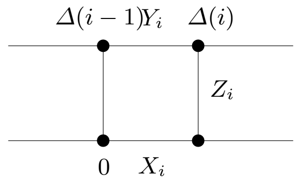

Fig. 1. Moving from Δ_{i-1} to Δ_i

path is the least constrained; denote this path $P^G_{\rm GP}(i)$. 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 $P^K_{\rm LR}(0,i)$ is the most expensive one under consideration. By the nature of the width-2 strip, the restriction of $P^G_{\rm GP}(i)$ 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 $\ell^G_{\rm GP}(i) \le \ell^K_{\rm LR}(0,i) \le \ell^G_{\rm GP}(i)+2$, and all the other lengths must also lie in this range. □

Because of Lemma 1, 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 $\ell^H_{\rm LR}(1,i)$ simply as $\ell(1,i)$ and $\ell^H_{\rm LR}(0,i)$ as $\ell(0,i)$ or just $\ell(i)$. Define

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

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 $i \ge 0$ let Z_i be the cost of $\{(0,i),(1,i)\}$. (See Figure 1 for visual reference.)

Observe that for $i > 0$,

$$\gamma(i) := \ell(i) - \ell(i-1) = \min\{X_i, \Delta(i-1) + Y_i + Z_i\} \quad (1)$$

$$\Delta(i) = \mathrm{trunc}(\Delta(i-1) + Y_i - X_i;\ -Z_i,\ Z_i). \quad (2)$$

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 $\mathrm{Be}(p)$, 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 typically not independent).

Lemma 2. *When the edge weights are i.i.d. with distribution* $\mathrm{Be}(p)$, $0 \le p < 1$, $\Delta(i)$ *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 $\{-1, 0, 1\}$.

By definition, the stationary distribution of the Markov chain is independent of its initial state, so we may assume that $\Delta(0) = 0$. In this case, the initial state is symmetric, and so is the transition rule, so the distribution of $\Delta(i)$ is symmetric for every i.

From (2), if $Z_i = 0$ (which occurs w.p. $1-p$) then $\Delta(i) = 0$. Otherwise we have the following table of possibilities, their probabilities (including the probability p that $Z_i = 1$, and defining $q := \mathbb{P}[\Delta(i-1) = 1] = \mathbb{P}[\Delta(i-1) = -1]$), 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.) □

Lemma 3. *When the edge weights are i.i.d. random variables with distribution* $\mathrm{Be}(p)$*,* $0 < p < 1$*,* $\gamma(i) = \ell(i) - \ell(i-1)$ *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 $p\left[qp^2 + (1-2q)(1-(1-p)^2) + q\right]$. The rest of the calculation is routine. □

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

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

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

5 Uniform Case: Expectation, Distribution, and Stationarity

What if X_i, Y_i, and Z_i have uniform distribution, $U[0,1]$? As in the previous sections, $\gamma(i)$ and $\Delta(i)$ are again Markov chains, but now with continuous state space. To avoid working with a continuous state space we will discretize it. First, for any (large) integer k, define $\underline{U}_k$ (resp., $\overline{U}_k$) to be the uniform distribution on the set $\{0, 1/k, \ldots, (k-1)/k\}$ (resp., $\{1/k, \ldots, (k-1)/k, 1\}$). Note that rounding a random variable $X \sim U[0,1]$ down and up to multiples of $1/k$ gives variables $\underline{X} \sim \underline{U}_k$ and $\overline{X} \sim \overline{U}_k$.

It is a simple observation that rounding all values X, Y and Z down or up gives (respectively) lower and upper bounds on any length $\ell(i)$. This allows bounds $\mathbb{E}\underline{\ell}_k(i) \leq \mathbb{E}\ell(i) \leq \mathbb{E}\overline{\ell}_k(i)$ to be computed much as in the Bernoulli case, via a finite Markov chain. By analogy with Theorem 4 and its proof (see full paper for details of the approach summarized in this paragraph), the first-passage percolation time constant can then be bounded by $\lim_{n\to\infty} \mathbb{E}[\underline{\ell}_k(n)/n] \leq \mathbb{E}[\gamma(n)] \leq \lim_{n\to\infty} \mathbb{E}[\overline{\ell}_k(n)/n]$. However, it is not true, for example, that $\gamma(i) \geq \underline{\ell}_k(i) - \underline{\ell}_k(i-1)$, and this natural approach thus characterizes γ's expectation but fails to say anything about its *distribution*. The distribution of γ may be of interest in itself, and that of Δ (which is essentially equivalent under (1)) is essential for computing quantities such as the expected VCG overpayment in the uniform model (paralleling its computation in the Bernoulli model in Section 6).

A different way of obtaining a finite Markov chain *does* provide us with random variables $\underline{\Delta}(n) \leq \Delta(n) \leq \overline{\Delta}(n)$, where $\underline{\Delta}(n)$ and $\overline{\Delta}(n)$ are given by finite Markov chains, allowing us to characterizes the distribution of $\Delta(n)$ and thereby giving access to any quantity of interest. Conceptually this method is quite different from the "make everything shorter / longer)" approach of the previous paragraph, but it is no harder: we simply derive what we want from the recurrence (2).

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

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

where $\underline{Z}$, $\overline{Z}$, *etc.* are any lower and upper bounds on their respective quantities. Specifically, taking $\underline{Z}$, $\overline{Z}$, and $\underline{W}$ to be the rounded-down and rounded-up discretizations of the respective variables, for any $\underline{\Delta} \leq \Delta$, and for convenience defining $\epsilon = 1/k$, we have

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

Thus,

$$\underline{\Delta}' := \underline{W} + \operatorname{trunc}(\underline{\Delta};\ -\underline{Z} - \underline{W} - \epsilon,\ \underline{Z} - \underline{W}) \tag{3}$$

ensures $\Delta' \geq \underline{\Delta}'$, and is thus a recursive formula for lower bounds. Similarly,

$$\overline{\Delta}' := \underline{W} + \epsilon + \operatorname{trunc}(\overline{\Delta};\ -\underline{Z} - \underline{W} - \epsilon,\ \underline{Z} - \underline{W}) \tag{4}$$

defines a recursion for upper bounds. As initial conditions we set $\underline{\Delta}(0) = -1$ and $\overline{\Delta}(1) = 1$ (deterministically), ensuring $\underline{\Delta}(0) \leq \Delta(0) \leq \overline{\Delta}(0)$, whereupon following equation (3) to define $\underline{\Delta}(n) = \underline{\Delta}'$ from $\underline{\Delta}(n-1) = \underline{\Delta}$ and likewise for equation (4) ensures that for all n, $\underline{\Delta}(n) \leq \Delta(n) \leq \overline{\Delta}(n)$. From (1), trivially,

$$\gamma(i) \geq \min\{\underline{X}_i, \underline{\Delta}(i-1) + \underline{Y} + \underline{Z}_i\} \tag{5}$$

$$\gamma(i) \leq \min\{\overline{X}_i, \overline{\Delta}(i-1) + \overline{Y} + \overline{Z}_i\} \tag{6}$$

Theorem 6 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 (5) and (6) defines corresponding random variables $\underline{\gamma}$ and $\overline{\gamma}$, which are then stochastic lower and upper bounds on $\gamma(n)$. Distribution functions for $\underline{\gamma}$ and $\overline{\gamma}$ are plotted in Figure 2. 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 with $k = 150$ we obtain $0.4215 < \mathbb{E}\gamma < 0.4292$. Computational aspects are discussed in the full paper.

This method allows us to get arbitrarily good estimates of the distribution of $\Delta(n)$, for n large (and thus, by Theorem 6, of the stationary distribution of Δ). It suffices to show that, for k large, the stationary random variables $\underline{\Delta}$ and $\overline{\Delta}$ are 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 $\epsilon = 1/k$.

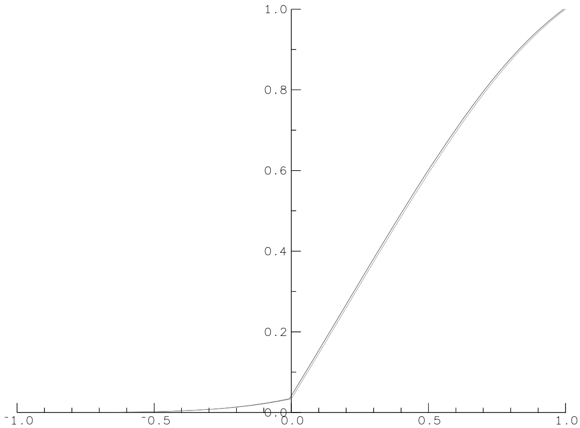

Fig. 2. Distribution functions for the stationary distributions of $\underline{\gamma}$ and $\overline{\gamma}$ given by $k = 150$. For any $n > 10$, the true incremental-length distribution $\gamma(n)$ lies between the two (incidentally proving that $0.4215 < \mathbb{E}\gamma < 0.4292$).

Theorem 5. *The stationary random variables* $\underline{\Delta}$ *and* $\overline{\Delta}$ *for equations* (3) *and* (4) *satisfy* $d(\underline{\Delta}, \overline{\Delta}) = O(\epsilon)$.

The proof is given in the full paper.

Arbitrarily good bounds on the stationary distribution of $\gamma(n)$, for n large, follow as a corollary. From (1), a random variable $\underline{\gamma}_i := \min\{\underline{X}_i, \underline{\Delta}(i-1)\}$ provides a lower bound $\underline{\gamma}_i \preceq \gamma_i$. Similarly, an upper bound is given 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, with $\overline{\Delta}$, $\underline{\Delta}$, $\overline{\gamma}$, and $\underline{\gamma}$ 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 unique well-defined stationary distribution; from (1) it is then immediate that γ does as well.

Theorem 6. *The continuous Markov chain* $\Delta(i)$ *defined by* (2) *has a unique stationary distribution.*

Proof. Per the remarks after Definition 7, any recurrent Harris chain possesses a unique stationary distribution, and Lemma 8 shows that $\Delta(i)$ is a recurrent Harris chain. □

Definition 7. *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.) It is known (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.

Lemma 8. *$\Delta(i)$ is a recurrent Harris chain, with $A = [-0.1, 1], B = [0, 0.4]$, and ν the uniform probability distribution on B multiplied by* 0.2.

Proof. 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 $\Delta(i+1) = \mathrm{trunc}(\Delta(i) + W_{i+1};\ -Z_{i+1},\ Z_{i+1}) = Z_{i+1}$. Let $V_{i+1} = \mathbf{1}\{Z_{i+1} \leq 0.4\}$. 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

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

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 7. 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 $\Delta(i) = \mathrm{trunc}(\Delta(i-1) + W_i;\ -Z_i,\ Z_i) \geq -0.1$, that is $\Delta(i) \in A$. Clearly this happens infinitely often w.p. 1. □

6 An Auction Model

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). 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 acting truthfully. The VCG mechanism finds an outcome that maximizes social welfare in a truthful fashion. For buying an (s, t)-path, the VCG mechanism is the following: An "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.)

Unfortunately, the VCG mechanism may result in the auctioneer paying much more than the cost of the shortest path. 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

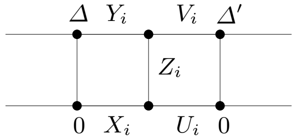

Fig. 3. The VCG cost at step i, working in from the left and the right, and assuming both sides are in stationarity

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 costs C and $C(1+\epsilon)$, and with the cheaper path containing k edges; the total payment is $C(1+k\epsilon)$.

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. (Average-case studies on other distributions of networks appear in [MPS03, CR04, KN05].)

Theorem 9. *When the edge weights are i.i.d.* $\mathrm{Be}(p)$ *random variables, with* $0<p<1$, *the VCG path cost satisfies*

$$\lim_{n\to\infty}\frac{1}{n}\mathbb{E}(\ell_{\mathrm{VCG}}(n)) = \frac{p\,(2+5\,p+4\,p^2+8\,p^3+11\,p^4-3\,p^6+p^8)}{(1+3\,p^2)^2}. \tag{7}$$

Proof. With reference to Figure 1, we compute the contribution of the ith 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 3, 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. $\mathrm{Be}(p)$ 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 (7).

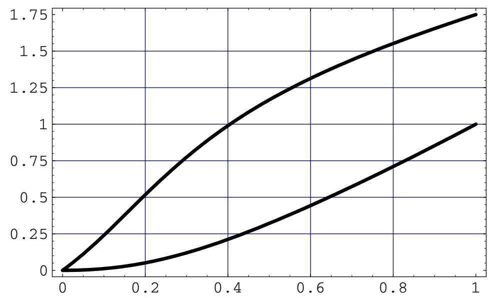

Fig. 4. VCG and usual shortest-path rates

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

$$C(X) = \mathbf{1}\{\min(A, B) \le \min(C, D)\} \cdot [\min(C, D) - (\min(A, B) - X)],$$

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

$$C(Y) = \mathbf{1}\{\min(C, D) < \min(A, B)\} \cdot [\min(A, B) - (\min(C, D) - Y)].$$

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

$$\begin{aligned} C(Z) = \mathbf{1}\{&(B < A) \vee (B \le \min(C, D)) \\ &\vee (D < \min(A, B, C))\} \cdot [\min(A, C) - (\min(B, D) - Z)]. \end{aligned}$$

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

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

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 (7). □

A plot of the VCG cost rate $\lim_{n\to\infty} \mathbb{E}\ell_{\mathrm{VCG}}(n)/n$, along with the corresponding shortest-path cost rate, is given in Figure 4.

References

[Ald01] David J. Aldous, *The $\zeta(2)$ limit in the random assignment problem*, Random Structures Algorithms **18** (2001), no. 4, 381–418.

[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.

[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.

[GNSar] D. Gamarnik, T. Nowicki, and G. Swirscsz, *Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method*, Random Structures and Algorithms (To appear).

[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.

[KN05] David Karger and Evdokia Nikolova, *Brief announcement: on the expected overpayment of VCG mechanisms in large networks*, PODC '05: Proceedings of the 24th annual ACM symposium on Principles of distributed computing (New York, NY, USA), ACM Press, 2005, pp. 126–126.

[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.