

1-1-2006

The Linking Probability of Deep Spider-Web Networks

Nicholas Pippenger
Harvey Mudd College

Recommended Citation

Pippenger, Nicholas. "The Linking Probability of Deep Spider-Web Networks." *SIAM Journal on Discrete Mathematics* 20, no. 1 (January 2006): 143-159.

This Article is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

THE LINKING PROBABILITY OF DEEP SPIDER-WEB NETWORKS*

NICHOLAS PIPPENGER†

Abstract. We consider crossbar switching networks with base b (that is, constructed from $b \times b$ crossbar switches), scale k (that is, with b^k inputs, b^k outputs, and b^k links between each consecutive pair of stages), and depth l (that is, with l stages). We assume that the crossbars are interconnected according to the spider-web pattern, whereby two diverging paths reconverge only after at least k stages. We assume that each vertex is independently idle with probability q , the vacancy probability. We assume that $b \geq 2$ and the vacancy probability q are fixed, and that k and $l = ck$ tend to infinity with ratio a fixed constant $c > 1$. We consider the linking probability Q (the probability that there exists at least one idle path between a given idle input and a given idle output). In a previous paper [*Discrete Appl. Math.*, 37/38 (1992), pp. 437–450] it was shown that if $c \leq 2$, then the linking probability Q tends to 0 if $0 < q < q_c$ (where $q_c = 1/b^{(c-1)/c}$ is the critical vacancy probability) and tends to $(1 - \xi)^2$ (where ξ is the unique solution of the equation $(1 - q(1 - x))^b = x$ in the range $0 < x < 1$) if $q_c < q < 1$. In this paper we extend this result to all rational $c > 1$. This is done by using generating functions and complex-variable techniques to estimate the second moments of various random variables involved in the analysis of the networks.

Key words. communication networks, crossbar switching networks, blocking probability

AMS subject classifications. 94C15, 60C05

DOI. 10.1137/050624376

1. Introduction. We deal in this paper with linking in crossbar switching networks, a phenomenon not dissimilar to that of percolation in lattices (as introduced by Broadbent and Hammersley [B] and surveyed by Grimmett [G]). An important difference, however, is that while percolation can be studied in finite subgraphs of a single infinite graph modeling the lattice, there is no single graph that naturally hosts the graph modeling crossbars switching networks in which we are interested. Our first order of business will be to describe these graphs.

A *crossbar graph* is characterized by three parameters: its *base*, $b \geq 2$, its *scale*, $k \geq 0$, and its *depth*, $l \geq 0$. Its vertices are partitioned into $l + 1$ *ranks*, each containing b^k vertices, which are labeled with the strings of length k over the alphabet $\{0, \dots, b - 1\}$. The vertices in rank 0 are called *inputs*, those in rank l are called *outputs*, and those in all other ranks are called *links*. The edges of the graph are partitioned into l *stages*, each containing b^{k+1} edges. For $1 \leq m \leq l$, the edges of stage m are directed out of vertices in rank $m - 1$ and into vertices in rank m . In a *spider-web* crossbar graph, which is our main concern in this paper, there is an edge of stage m from vertex v of rank $m - 1$ to vertex w of rank m if and only if v and w are labeled by strings that differ at most in position j , where $j \equiv m \pmod{k}$. The edges of each stage are thus partitioned into b^{k-1} $b \times b$ complete bipartite graphs (called *crossbars*). The spider-web crossbar graph with base b , scale k , and depth l will be denoted $G_{b,k,l}$. We shall see in section 2 that if $l \geq k$, there are b^{l-k} paths from a given input to a given output; if $l < k$, there is at most one path from a given input to a given output. Our main interest is in spider-web crossbar graphs with $l \geq k$, since

*Received by the editors February 14, 2005; accepted for publication (in revised form) September 15, 2005; published electronically March 3, 2006. This work was supported in part by NSF grant CCF-0430656.

<http://www.siam.org/journals/sidma/20-1/62437.html>

†Department of Computer Science, Princeton University, 35 Olden Street, Princeton, NJ 08540 (njp@princeton.edu).

in these graphs any input can be connected by a path to any output; in our analysis, however, graphs with $l < k$ will occur as subgraphs, so it will be necessary to allow this case in some of our intermediate results.

We shall assume that each vertex in the graph $G_{b,k,l}$ is independently assigned the status *idle*, with probability q (called the *vacancy* probability), or *busy*, with the complementary probability $p = 1 - q$ (called the *occupancy* probability). This random assignment of a status to each vertex in a graph will be called the *state* of the graph. Given an input v and output w , let $Q_{v,w}$ (called the *linking* probability) denote the probability that there exists a path consisting entirely of idle links from v to w . (In this paper, “path” will always mean “directed path.” In general, the linking probability is defined as the *conditional* probability that there exists an idle path, *given* that v and w are themselves idle, but for the probabilistic model that we are using, this condition is independent.) We shall see in section 2 that if $l \geq k$, the probability $Q_{v,w}$ does not depend on the choice of the input-output pair (v, w) , so we shall let Q denote the common value of these probabilities. The complementary probability $P = 1 - Q$ (called the *blocking* probability) is the probability that all paths between a given input-output pair (v, w) are broken by a set of busy links.

In practice, the parameter p represents the amount of traffic being carried by a crossbar network (which one would like to maximize), and the parameter P represents the fraction of arriving traffic lost due to congestion within the network (which one would like to minimize). In analysis, however, it is almost always more convenient to work with the complementary parameters q and Q , so we shall work exclusively with these parameters in what follows.

In practice, a graph $G_{b,k,l}$ would be fixed, and the linking probability Q would be studied as a function of the vacancy probability q . It is found that Q undergoes a rapid transition from a value near zero to a significantly positive value as q passes through a neighborhood of $1/b^{(l-k)/(l-1)}$. This is easily understood in the following way.

Let the random variable $X_{v,w}$ denote the number of idle paths from v to w . We shall see in section 2 that if $l \geq k$, the distribution of $X_{v,w}$ does not depend on the choice of the input-output pair (v, w) , so we shall let X denote a random variable with this common distribution. Each of the b^{l-k} paths from v to w contains $l - 1$ links, which are all idle with probability q^{l-1} . Thus we have

$$(1.1) \quad \text{Ex}[X] = b^{l-k} q^{l-1}.$$

Thus as q passes through $1/b^{(l-k)/(l-1)}$, the expected number of idle paths from v to w (called the *specific transparency*) goes from an exponentially decreasing to an exponentially increasing function of k and l . This suggests that if k and l tend to infinity in such a way that their ratio $c = l/k > 1$ remains fixed while b and q are also held fixed, then Q will tend to a limit, and this limit will have a discontinuity as q passes through the critical value

$$q_c = 1/b^{(c-1)/c}.$$

(We note that $1 < c < \infty$ implies $1/b < q_c < 1$.) Our goals in this paper are to confirm this conjecture and to determine the limiting value of Q .

Our first step toward these goals, taken in section 2, will be to derive the following estimate for the second moment $\text{Ex}[X^2]$ of X .

THEOREM 1.1. *Let both $b \geq 2$ and $1/b < q < 1$ be fixed. Then*

$$\mathbb{E}x[X^2] = \mathbb{E}x[X] \cdot \left(\left(\frac{b-1}{bq-1} \right)^2 b^{l-k} q^{l+1} + 1 + O(l b^{l-2k} q^l) + O(l q^k) \right)$$

as $k, l \rightarrow \infty$ with $l \geq k$ and $(\log(l+1))/(k+1) \rightarrow 0$. (The constants in the O -terms may depend on b and q , but are independent of k and l .)

We observe that this estimate is enough to establish that the limiting value (if it exists) of Q for $k \rightarrow \infty$ and $l = ck$ cannot be a continuous function of Q as q passes through q_c . Indeed, from Markov's inequality and (1.1), we have

$$(1.2) \quad Q = \Pr[X \geq 1] \leq \mathbb{E}x[X] = b^{l-k} q^{l-1} \rightarrow 0$$

for $q < q_c$. On the other hand, (1.1) and Theorem 1.1, together with the inequality

$$(1.3) \quad \Pr[X \geq 1] \geq \frac{\mathbb{E}x[X]^2}{\mathbb{E}x[X^2]},$$

imply

$$(1.4) \quad \begin{aligned} Q = \Pr[X \geq 1] &\geq \frac{\mathbb{E}x[X]^2}{\mathbb{E}x[X^2]} = \frac{(bq-1)^2}{(b-1)^2 q^2 + (bq-1)^2 q} \left(1 + O\left(\frac{l}{b^k}\right) + O(l q^k) \right) \\ &\rightarrow \frac{(bq-1)^2}{(b-1)^2 q^2 + (bq-1)^2 q} > 0 \end{aligned}$$

for $q = q_c$. (To verify (1.3), we consider the distribution of X conditioned on the event $X \geq 1$. Since x^2 is a convex function of x , we have

$$\mathbb{E}x[X^2 \mid X \geq 1] \geq \mathbb{E}x[X \mid X \geq 1]^2.$$

Multiplying by $\Pr[X \geq 1]^2$ yields

$$\begin{aligned} \mathbb{E}x[X^2] \Pr[X \geq 1] &= \mathbb{E}x[X^2 \mid X \geq 1] \Pr[X \geq 1]^2 \\ &\geq \mathbb{E}x[X \mid X \geq 1]^2 \Pr[X \geq 1]^2 \\ &= \mathbb{E}x[X]^2, \end{aligned}$$

which is equivalent to (1.3).) The inequalities in (1.2) and (1.4) show that the inferior limit of Q for $q = q_c$ is strictly greater than the limiting value for $q < q_c$, as claimed.

The argument of the preceding paragraph also sheds some light on the condition $(\log(l+1))/(k+1) \rightarrow 0$ in Theorem 1.1. (This condition involves $k+1$ and $l+1$ rather than k and l simply to avoid dividing by or taking the logarithm of 0.) This condition is not the weakest one sufficient to give an estimate of the form $\mathbb{E}x[X^2] = O(\mathbb{E}x[X]^2)$, but it is clear that some upper bound on the growth of l must be imposed, for with probability $(1-q)^{b^k}$ all the links in a given rank are busy, disconnecting all input-output pairs. Thus if $l \cdot (1-q)^{b^k} \rightarrow \infty$, we have $Q \rightarrow 0$, contradicting the implication of (1.3) when $\mathbb{E}x[X^2] = O(\mathbb{E}x[X]^2)$.

In section 3, we shall combine Theorem 1.1 with branching-process arguments from Pippenger [P3] to establish the existence and determine the limiting value of Q for $q > q_c$.

THEOREM 1.2. *Let $b \geq 2$ and $0 < q < 1$ be fixed, and let $c > 1$ be rational. Then as $k \rightarrow \infty$ with $l = ck$, we have*

$$Q \rightarrow \begin{cases} 0 & \text{if } 0 < q < q_c, \\ (1 - \xi)^2 & \text{if } q_c < q < 1, \end{cases}$$

where ξ is the unique solution of the equation $x = (1 - q(1 - x))^b$ in the range $0 < x < 1$.

A comment is in order concerning the behavior of $(1 - \xi)^2$ as a function of q . The function $f(x) = (1 - q(1 - x))^b$ is a strictly convex function of x for $0 < q \leq 1$, since $f''(x) = b(b - 1)q^2(1 - q(1 - x))^{b-2} > 0$ in this range. Thus the graph of $f(x)$ can intersect the diagonal at most twice in this range. There is one intersection at $x = 1$, and the conditions $f(0) = (1 - q)^b > 0$ and $f'(1) = bq > 1$ imply that there is at least one intersection in the range $0 < x < 1$ when $1/b < q < 1$. Thus there is indeed a unique solution of the equation $x = (1 - q(1 - x))^b$ in the range $0 < x < 1$ when $1/b < q < 1$, and this latter condition is implied by $q_c < q < 1$. The degree of this equation can be reduced by 1 (because of the solution $x = 1$), and it is easy to see that the resulting equation is irreducible over the field of rational functions of q ; thus ξ is an algebraic function of q of degree $b - 1$. Since $(1 - \xi)^2$ is a polynomial in ξ , it is also an algebraic function of q of degree $b - 1$. Straightforward analysis shows that $Q \rightarrow 1$ as $q \rightarrow 1$ with $1 - Q = 1 - (1 - \xi)^2 \sim 1 - 2(1 - q)^b$, which may be interpreted as saying that the main obstacle to linking when $q \rightarrow 1$ is complete occupation either of the b links adjacent to the input in rank 1, or of the b links adjacent to the output in rank $l - 1$. As $q \rightarrow 1/b$ from above, we have $(1 - \xi)^2 \sim (bq - 1)^2 / \binom{b}{2}$.

Theorem 1.2 was proved, under the additional restriction $c \leq 2$, by Pippenger [P3], so the additional contribution of the current paper consists of lifting this restriction. Nevertheless, the techniques used in the current paper go considerably beyond those employed in the previous paper in that the proof of Theorem 1.1 starts with a detailed combinatorial examination of the intersections between paths, then uses complex-variable techniques to determine the asymptotics of the quantities involved.

Spider-web networks were introduced by Ikeno [I] (though the term *spider-web* has sometimes been used to refer to a broader class of networks). They have several optimality properties among networks constructed from the same type and number of crossbars. Takagi [T] showed that they have the largest linking probability in a large class of crossbar networks called “rhyming” networks. Chung and Hwang [C] showed that, surprisingly, these networks are *not* optimal in the larger class of “balanced” networks. But Pippenger [P3] showed that they are *asymptotically* optimal in this class for $1 < c \leq 2$, and the current paper extends this result to all $c > 1$.

The probability distribution on states that we use was introduced by Lee [L1] and Le Gall [L2, L3]. It is by far the easiest to use for analytical purposes, but it suffers from the defect that the set of busy vertices does not form a set of coherent paths from inputs to outputs. Models addressing this defect have been introduced by Koverninskii [K] and Pippenger [P1], and the results in [P3] have been extended to these models in [P2]. It seems likely that the results of the the present paper can be similarly extended.

The current paper is self-contained, except for some estimates concerning branching processes taken from Pippenger [P3]. We have followed the notation of that paper, except that the base, which was denoted d in that paper, is now denoted b (to free

the symbol d for its traditional use in the calculus).

2. The second moment. Our goal in this section is to prove Theorem 1.1. We begin with a combinatorial result concerning spider-web graphs.

LEMMA 2.1. *The automorphism group of $G_{b,k,l}$ acts transitively on the paths from inputs to outputs.*

Proof. Since an automorphism must permute the vertices within each rank, an automorphism ϑ may be regarded as a sequence $\vartheta = (\vartheta_0, \dots, \vartheta_l)$ of permutations, one for each rank. We shall focus on automorphisms in which each ϑ_m (for $0 \leq m \leq l$) is characterized by a string $\vartheta_{m,1} \cdots \vartheta_{m,k}$ of k digits from the alphabet $\{0, \dots, b-1\}$ and acts on the vertices of rank m by carrying the vertex labeled $a_1 \cdots a_k$ to the vertex labeled $a'_1 \cdots a'_k$, where $a'_j \equiv a_j + \vartheta_{m,j} \pmod{b}$ for $1 \leq j \leq k$. If, for $1 \leq m \leq l$, the string ϑ_{m-1} differs from the string ϑ_m in at most position j , where $j \equiv m \pmod{k}$, then the sequence $\vartheta = (\vartheta_0, \dots, \vartheta_l)$ will constitute an automorphism.

To show that the automorphisms act transitively on the paths, it will suffice to show, for some fixed path u^* , that for every path u , there is an automorphism that carries u^* to u (since then the inverse of such an automorphism can be used to carry any other path u' to u^*). A path u may be regarded as a sequence $u = (u_0, \dots, u_l)$ of vertex labels in which, for $1 \leq m \leq l$, the string u_{m-1} differs from the string u_m in at most position j , where $j \equiv m \pmod{k}$. We shall choose for u^* the path $u^* = (0^k, \dots, 0^k)$. Then clearly the automorphism $\vartheta = (\vartheta_0, \dots, \vartheta_l)$ defined by $\vartheta_m = u_m$ for $0 \leq m \leq l$ carries u^* to u . \square

COROLLARY 2.2. *If $l \geq k$, the graph $G_{b,k,l}$ contains b^{l-k} paths from any given input to any given output; if $l < k$, there is at most one path from any given input to any given output.*

Proof. If $l \geq k$, every input-output pair is joined by at least one path, since every position in the strings labeling vertices has an opportunity to change at least once. Thus, by Lemma 2.1 every input is joined by the same number of paths. Since each of the b^k inputs is the origin of b^l paths to outputs, there are a total of b^{l+k} paths joining inputs to outputs, and thus b^{l-k} paths joining each of the b^{2k} input-output pairs. If $l < k$, there is a path from input v to output w only if the labels of v and w agree in the last $k-l$ positions. Thus $G_{b,k,l}$ breaks into b^{k-l} disjoint components, each containing b^l vertices in each rank; there is a unique path joining input v to output w if they belong to the same component, but no path joining them if they belong to different components. \square

COROLLARY 2.3. *If $l \geq k$, the automorphism group of $G_{b,k,l}$ acts transitively on the input-output pairs.*

Proof. If $k \geq k$, each input-output pair is joined by a path, so the corollary follows from Lemma 2.1. \square

This corollary, together with the fact that the probability distribution on states of the graph is invariant under automorphisms of the graph, justifies our earlier assertion that the linking probability $Q_{v,w}$ and the distribution of the random variable $X_{v,w}$ are independent of the choice of the input-output pair (v, w) when $l \geq k$. Henceforth we shall focus our attention on the input-output pair $(v^*, w^*) = (0^k, 0^k)$. If $l \geq k$, this entails no loss of generality. When $l < k$, we shall deal only with cases in which the input and output of interest are joined by a path, and in these cases there is again no loss of generality.

Fix $b \geq 2$ and $k \geq 1$. For $l \geq 0$, let $\varphi_l(y)$ denote the generating function for the number of paths from the input $v^* = 0^k$ to the output $w^* = 0^k$ classified according to the number of links that have labels different from 0^k ; that is, the coefficient of y^m in

$\varphi_l(y)$ is the number of paths from v^* to w^* that have $l-1-m$ links in common with the path $u^* = (0^k, \dots, 0^k)$. Clearly $\varphi_l(y) = 1$ for $0 \leq l \leq k$, and $\varphi_l(y)$ is a polynomial in y of degree $l-1$ if $l \geq k+1$.

We are interested in the polynomials $\varphi_l(y)$ for various values of $l \geq 0$, with b and k fixed. To determine them, it will be convenient to work with a graph $G_{b,k}$ that contains as subgraphs all the graphs $G_{b,k,l}$ for various values of l . For any $m \geq l \geq 0$, $G_{b,k,l}$ may be regarded as the subgraph comprising the vertices in ranks 0 through l and the edges in stages 1 through l of $G_{b,k,m}$. Thus we may define the infinite graph

$$G_{b,k} = \bigcup_{l \geq k} G_{b,k,l}$$

as the union (inductive limit) of all these graphs. The graph $G_{b,k}$ has inputs in rank 0, but all other vertices will be referred to as links.

For $l \geq 0$, the polynomial $\varphi_l(y)$ is the generating function for the number of paths from the input $v^* = 0^k$ to the link labeled 0^k in rank l classified according to the number of links that have labels different from 0^k .

Let

$$\psi(y, z) = \sum_{l \geq 0} \varphi_l(y) z^l$$

be the generating function for the polynomials $\varphi_l(y)$. The key to our estimate for the second moment of X is the following proposition.

PROPOSITION 2.4. *We have*

$$\psi(y, z) = \frac{1 - byz + (b-1)(yz)^{k+1}}{(1-z)(1-byz) - (b-1)z(1-y)(yz)^k}.$$

Proof. In this proof, we shall employ a concise alternative representation of a path $u = (u_0, \dots, u_l)$ of length $l \geq 0$ as a string $t = t_1 \dots t_{k+l}$ of length $k+l$ over the alphabet $B = \{0, \dots, b-1\}$. The first k digits $t_1 \dots t_k$ of t will be the k digits of the label u_0 . For $1 \leq m \leq l$, t_{k+m} will be the digit in position j of u_m , where $j \equiv m \pmod{k}$ (the digit of u_m that might be different from that of u_{m-1}). Then for $0 \leq m \leq l$, u_m is the string $t_{m+1} \dots t_{m+k}$. In particular, the last k digits of t are the k digits of the label u_l of the link in rank l , and the paths from the input $v^* = 0^k$ to the link labeled 0^k in rank l are in one-to-one correspondence with the strings of length $k+l$ over the alphabet B , whose first k digits and last k digits are 0's.

Given a path $t = 0^k t_{k+1} \dots t_{l-k} 0^k$, let us overline each digit t_{k+m} ($1 \leq m \leq l$) for which $u_{m-1} \neq 0^k$. The result is a string over the alphabet $B \cup \overline{B}$, where $\overline{B} = \{\overline{0}, \dots, \overline{b-1}\}$ is the set of overlined digits. For $l \geq 0$, let the language $K_l \subseteq (B \cup \overline{B})^{k+l}$ comprise the strings obtained in this way for all paths from the input $v^* = 0^k$ to the link labeled 0^k in rank l , and define $K \subseteq (B \cup \overline{B})^*$ by

$$K = \bigcup_{l \geq 0} K_l.$$

Then $\psi(y, z)$ is the power series in y and z in which the coefficient of $y^j z^l$ is the number of strings of length $k+l$ in K in which j digits are overlined. Let

$$L = 0^{-k} K = \{t \in (B \cup \overline{B})^* : 0^k t \in K\}$$

be the language obtained from K by deleting the k initial 0's from each string. Since none of this initial 0's are overlined, $\psi(y, z)$ is the power series in y and z in which the coefficient of $y^j z^l$ is the number of strings of length l in L in which j digits are overlined.

Our next step is to write a regular expression for the language L . Define the alphabets $B' = \{1, \dots, b-1\}$ and $\overline{B}' = \{\overline{1}, \dots, \overline{b-1}\}$. Then L is described by the regular expression

$$(2.1) \quad \left(\left(\Lambda + \left(\overline{B}' \left(\Lambda + \overline{0} + \dots + \overline{0}^{k-1} \right) \right)^* \overline{B}' \overline{0}^{k-1} \right) 0 \right)^*,$$

where Λ denotes the empty string. To see this, we observe that a string in L can be uniquely parsed into zero or more *stretches*, each of which ends with an unoverlined 0. A stretch consists of an unoverlined 0 optionally preceded by an *excursion*. An excursion consists of a *final segment* preceded by zero or more *preliminary segments*. A final segment consists of a digit from \overline{B}' followed by exactly $k-1$ overlined 0's. A preliminary segment consists of a digit from \overline{B}' followed by at most $k-1$ overlined 0's. Clearly a final segment is described by the regular expression $\overline{B}' \overline{0}^{k-1}$, and a preliminary segment is described by the regular expression $\overline{B}' (\Lambda + \overline{0} + \dots + \overline{0}^{k-1})$. Thus an excursion is described by the regular expression

$$\left(\overline{B}' \left(\Lambda + \overline{0} + \dots + \overline{0}^{k-1} \right) \right)^* \overline{B}' \overline{0}^{k-1},$$

and a stretch is described by the regular expression

$$\left(\Lambda + \left(\overline{B}' \left(\Lambda + \overline{0} + \dots + \overline{0}^{k-1} \right) \right)^* \overline{B}' \overline{0}^{k-1} \right) 0.$$

Thus the strings in L are described by the regular expression (2.1).

We now observe that the regular expression (2.1) is *unambiguous* in the following sense: A string described by a subexpression $R + S$ is described by R or by S (but not both); a string t described by a subexpression RS has a unique parsing $t = rs$ such that r is described by R and s is described by S ; and a string t described by a subexpression S^* has a unique parsing $s = s_1 \dots s_n$ with $n \geq 0$ such that s_1, \dots, s_n are described by S .

For an unambiguous regular expression, if $\psi_R(y, z)$ and $\psi_S(y, z)$ are the generating functions counting the strings described by subexpressions R and S , respectively, then $\psi_R(y, z) + \psi_S(y, z)$, $\psi_R(y, z) \psi_S(y, z)$, and $1/(1 - \psi_S(y, z))$ are the generating functions counting the strings described by the subexpressions $R + S$, RS , and S^* , respectively.

Thus the final segments are counted by the generating function $(b-1)(yz)^k$ and the preliminary segments are counted by the generating function

$$(b-1)yz(1 + yz + \dots + (yz)^{k-1}) = \frac{(b-1)(yz - (yz)^{k+1})}{1 - yz}.$$

The excursions are counted by

$$\frac{(b-1)(yz)^k}{1 - \frac{(b-1)(yz - (yz)^{k+1})}{1 - yz}} = \frac{(b-1)((yz)^k - (yz)^{k+1})}{1 - yz - (b-1)(yz - (yz)^{k+1})},$$

and the stretches are counted by

$$\left(1 + \frac{(b-1)((yz)^k - (yz)^{k+1})}{1 - yz - (b-1)(yz - (yz)^{k+1})} \right) z = \frac{z - yz^2 - (b-1)z(yz - (yz)^k)}{1 - yz - (b-1)(yz - (yz)^{k+1})}.$$

Thus the strings in L are counted by

$$\frac{1}{1 - \frac{z-yz^2-(b-1)z(yz-(yz)^k)}{1-yz-(b-1)(yz-(yz)^{k+1})}} = \frac{1 - byz + (b-1)(yz)^{k+1}}{(1-z)(1-byz) - (b-1)z(1-y)(yz)^k},$$

which completes the proof of the proposition. \square

PROPOSITION 2.5. *Let $b \geq 2$ and $0 < q < 1$ be fixed. Then as $k \rightarrow \infty$, and as $l \geq 0$ behaves in such a way that $(\log(l+1))/(k+1) \rightarrow 0$, we have*

$$\varphi_l(q) = \left(\frac{b-1}{bq-1}\right)^2 b^{l-k} q^{l+1} + 1 + O(lb^{l-2k} q^l) + O(lq^k) + O(lq^l).$$

(The constants in the O -terms may depend on b and q , but are independent of k and l .)

Proof. Write $A(z) = 1 - bqz + (b-1)(qz)^{k+1}$ and $B(z) = (1-z)(1-bqz) - (b-1)z(1-q)(qz)^k$ so that $\psi(q, z) = A(z)/B(z)$. Then from Cauchy's formula we have

$$\begin{aligned} \varphi_l(q) &= \frac{1}{2\pi i} \oint_{\Gamma_0} \frac{\psi(q, z) dz}{z^{l+1}} \\ (2.2) \qquad &= \frac{1}{2\pi i} \oint_{\Gamma_0} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}}, \end{aligned}$$

where Γ_0 is a contour taken counterclockwise around a circle $|z| = \varepsilon$ centered at 0 and having radius ε sufficiently small to exclude all other singularities of the integrand.

To make further progress, we must estimate the locations of these other singularities, which are poles at the values of z for which the denominator $B(z)$ vanishes. One such singularity is at $z = 1/q$. Let

$$\zeta_1 = \frac{1}{q} \left(1 - \frac{1}{l}\right),$$

and let Γ_1 be a contour taken counterclockwise around the circle $|z| = \zeta_1$ centered at 0 and having radius ζ_1 . As z traverses this contour, the magnitude of the first term $(1-z)(1-bqz)$ of $B(z)$ satisfies the lower bound

$$\begin{aligned} |(1-z)(1-bqz)| &= |1-z| \cdot |1-bqz| \\ &\geq \left(\frac{1}{q} - 1 - \frac{1}{ql}\right) \left(b - 1 - \frac{b}{l}\right) \\ &\geq \left(\frac{1}{q} - 1\right) (b-1) - \frac{bq-1}{ql}, \end{aligned}$$

since the minimum occurs when z is real and positive. The magnitude of the second term, $(b-1)z(1-q)(qz)^k$, on the other hand, satisfies the upper bound

$$\begin{aligned} |(b-1)z(1-q)(qz)^k| &\leq (b-1) \left(\frac{1}{q} - 1\right) \left(1 - \frac{1}{l}\right)^{k+1} \\ &\leq (b-1) \left(\frac{1}{q} - 1\right) e^{-k/l} \\ &\leq (b-1) \left(\frac{1}{q} - 1\right) \left(1 - \frac{(e-1)k}{el}\right). \end{aligned}$$

Here we have used the inequality $1-x \leq e^{-x}$, which holds for all x because the graph of the convex function e^{-x} lies above that of $1-x$, its tangent at $x = 0$, and the inequality $e^{-x} \leq 1 - (e-1)x/e$, which holds for $0 \leq x \leq 1$ because the graph of the convex function e^{-x} lies below that of $1 - (e-1)x/e$, its chord across the interval $0 \leq x \leq 1$. Thus for all sufficiently large k (specifically, for $k > (bq-1)e/(b-1)(1-q)(e-1)$), we have the bound

$$|B(z)| = \Omega\left(\frac{1}{l}\right)$$

for z on the contour Γ_1 . Since we also have $A(z) = O(1)$ for z on Γ_1 , we have the estimate

$$(2.3) \quad \frac{1}{2\pi i} \oint_{\Gamma_1} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} = O(lq^l).$$

Furthermore, as z traverses the contour Γ_1 , the value of the first term, $(1-z)(1-bqz)$, in $B(z)$ circles the origin twice, since it is a quadratic polynomial. Since the second term, $(b-1)z(1-q)(qz)^k$, has strictly smaller magnitude, the value of $B(z)$ also circles the origin twice. It follows that the denominator of $B(z)$ has exactly two zeros inside the contour Γ_1 . These are perturbations of the zeros of the first term: the zero of the first term at $z = 1$ is perturbed to one at

$$(2.4) \quad z = \zeta_2 = 1 + O(q^k),$$

and the zero of the first term at $z = 1/bq$ is perturbed to one at

$$(2.5) \quad z = \zeta_3 = \frac{1}{bq} \left(1 - \frac{(b-1)(1-q)}{(bq-1)b^k} + O\left(\frac{k}{b^{2k}}\right) \right).$$

The condition $(\log(l+1))/(k+1) \rightarrow 0$ ensures that the O -terms in (2.4) and (2.5) have smaller orders of magnitude than the terms preceding them. We observe that $0 < \zeta_3 < \zeta_2 < \zeta_1$, and thus $0, \zeta_3$, and ζ_2 lie inside Γ_1 and lie in that order along the real axis. Let Γ_2 be a contour taken counterclockwise around a circle $|z - \zeta_2| = \varepsilon$ centered at ζ_2 and having radius ε sufficiently small to exclude all other singularities of the integrand, and let Γ_3 be a contour taken counterclockwise around a circle $|z - \zeta_3| = \varepsilon$ centered at ζ_3 and having radius ε sufficiently small to exclude all other singularities of the integrand. Since the integral of an analytic function around a contour depends only on the homology class of the contour in the domain of analyticity of the function, and since Γ_0 is homologous to $\Gamma_1 - \Gamma_2 - \Gamma_3$ (indeed, Γ_0 is homotopic to a contour that joins a forward traversal of Γ_1 with reverse traversals of Γ_2 and Γ_3 by canceling traversals of segments $[\zeta_3 + \varepsilon, \zeta_2 - \varepsilon]$ and $[\zeta_2 + \varepsilon, \zeta_1]$ of the real axis), from (2.2) we have

$$(2.6) \quad \begin{aligned} \varphi_l(q) &= \frac{1}{2\pi i} \oint_{\Gamma_1} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} \\ &\quad - \frac{1}{2\pi i} \oint_{\Gamma_2} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} \\ &\quad - \frac{1}{2\pi i} \oint_{\Gamma_3} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}}. \end{aligned}$$

The first integral in (2.6) has already been estimated in (2.3). The remaining integrals circle just one singularity of the integrand, and thus they can be evaluated

by Cauchy's formula. If ζ is a simple pole of the integrand, and if Γ is a contour taken clockwise around just this singularity of the integrand, then we have

$$\begin{aligned} \frac{1}{2\pi i} \oint_{\Gamma} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} &= \operatorname{Res}_{z=\zeta} \frac{A(z)}{B(z)} \frac{1}{z^{l+1}} \\ &= \frac{A(\zeta)}{B'(\zeta)} \frac{1}{\zeta^{l+1}}. \end{aligned}$$

For the integral around Γ_2 , we have $A(\zeta_2) = -(bq - 1) + O(q^k)$ and $B'(\zeta_2) = (bq - 1) + O(kq^k)$ so that

$$(2.7) \quad -\frac{1}{2\pi i} \oint_{\Gamma_2} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} = 1 + O(lq^k).$$

For the integral around Γ_3 we have $A(\zeta_3) = (b - 1)^2(bq - 1)b^{k+1} + O(k/b^{2k})$ and $B'(\zeta_3) = -(bq - 1) + O(k/b^k)$ so that

$$(2.8) \quad -\frac{1}{2\pi i} \oint_{\Gamma_3} \frac{A(z)}{B(z)} \frac{dz}{z^{l+1}} = \left(\frac{b-1}{bq-1}\right)^2 b^{l-k} q^{l+1} + O(lb^{l-2k} q^l).$$

Substituting the estimates (2.3), (2.7), and (2.8) into (2.6) completes the proof of the proposition. \square

We observe that by extending the asymptotic expansions in (2.4) and (2.5), it is possible to extend the expansions in (2.7) and (2.8) and thus reduce their contributions to the error terms in Proposition 2.4. The error term in (2.3), however, cannot be improved without taking account of the zeros of $B(z)$ outside the circle $|z| = 1/q$, which will in general contribute oscillatory terms to the expansion of $\varphi_l(q)$.

Proof of Theorem 1.1. By Corollary 2.3, we may take X to be the number of idle paths from $v^* = 0^k$ to $w^* = 0^k$. We then have

$$(2.9) \quad \begin{aligned} \operatorname{Ex}[X^2] &= \sum_{u':v^* \rightarrow w^*} \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[u \text{ idle}, u' \text{ idle}] \\ &= \sum_{u':v^* \rightarrow w^*} \operatorname{Pr}[u' \text{ idle}] \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[u \text{ idle} \mid u' \text{ idle}], \end{aligned}$$

where the sums are over all paths from v^* to w^* . For each path u' , we can find by Lemma 2.1 an automorphism ϑ that carries u' to the path u^* in which all links are labeled 0^* . Applying this automorphism to both u and u' gives $\operatorname{Pr}[u \text{ idle} \mid u' \text{ idle}] = \operatorname{Pr}[\vartheta(u) \text{ idle} \mid u^* \text{ idle}]$, since the probability distribution on states of the graph is invariant under automorphisms. Furthermore,

$$\begin{aligned} \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[u \text{ idle} \mid u' \text{ idle}] &= \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[\vartheta(u) \text{ idle} \mid u^* \text{ idle}] \\ &= \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[u \text{ idle} \mid u^* \text{ idle}], \end{aligned}$$

since both right-hand sides sum the same terms in different orders. Thus the inner sum in (2.9) is independent of u' , and we have

$$\operatorname{Ex}[X^2] = \sum_{u':v^* \rightarrow w^*} \operatorname{Pr}[u' \text{ idle}] \sum_{u:v^* \rightarrow w^*} \operatorname{Pr}[u \text{ idle} \mid u^* \text{ idle}]$$

so that $\text{Ex}[X^2]$ factors as the product of two sums. The first sum is just $\text{Ex}[X]$. To evaluate the second sum, we observe that $\Pr[u \text{ idle} \mid u^* \text{ idle}]$ is just q^j , where j is the number of links on u that are not labeled 0^k . Thus the second sum is $\varphi_l(q)$, and the theorem follows from Proposition 2.5. \square

3. The linking probability. Our goal in this section is to prove Theorem 1.2. Thus in this section we shall always assume that $b \geq 2$ and $0 < q < 1$ are fixed and that $k \rightarrow \infty$ and $l = ck$ for some fixed rational $c > 1$. Thus the constants in O -terms may depend on c as well as on b and q , but not on k or l . We shall also assume that k is even; the case of odd k requires only that $k/2$ be replaced with $\lfloor k/2 \rfloor$ and $\lceil k/2 \rceil$ in appropriate ways.

LEMMA 3.1. *Let $G_{b,k,l}^*$ be the graph obtained from $G_{b,k,l}$ by reversing the direction of its edges and exchanging the roles of its inputs and outputs. Then $G_{b,k,l}^*$ is isomorphic to $G_{b,k,l}$.*

Proof. The isomorphism takes the vertex with label $a_1 \cdots a_k$ in rank m of $G_{b,k,l}$ to the vertex with label $a_1^* \cdots a_k^*$ in rank $l - m$ of $G_{b,k,l}^*$, where $a_i^* = a_j$ with $j \equiv l + 1 - i \pmod{k}$ (and, conversely, as it is an involution). \square

Lemma 3.1 establishes a symmetry between $G_{b,k,l}$ and $G_{b,k,l}^*$, which we shall invoke by use of the term “dually.” (When l is even, $G_{b,k,l}$ is in fact isomorphic to a graph with manifest bilateral symmetry, as is shown in the appendix of Pippenger [P3].)

LEMMA 3.2. *Let $\langle G_{b,k,l} \rangle_{m,n}$, with $0 \leq m \leq n \leq l$, be the subgraph of $G_{b,k,l}$ comprising the vertices in ranks m (now considered inputs) through n (now considered outputs) and the edges in stages $m + 1$ through n . Then $\langle G_{b,k,l} \rangle_{m,n}$ is isomorphic to $G_{b,k,n-m}$.*

Proof. The isomorphism takes the vertex with label $a_1 \cdots a_k$ in rank h of $G_{b,k,n-m}$ to the vertex with label $a'_1 \cdots a'_k$ in rank $m + h$ of $\langle G_{b,k,l} \rangle_{m,n}$, where $a'_i = a_j$ with $j \equiv i + m \pmod{k}$. \square

COROLLARY 3.3. *Between any given input and any given output of $\langle G_{b,k,l} \rangle_{m,n}$, there are b^{n-m-k} paths if $n - m \geq k$, and there is either one path or none if $n - m < k$.*

Proof. The proof is immediate from Lemma 3.2 and Corollary 2.2. \square

We begin with the upper bound to Q . For $0 < q < q_c$, where $q_c = 1/b^{(c-1)/c}$, we have $Q \rightarrow 0$ by (1.2). For $q_c < q < 1$, we shall use the following lemma from Pippenger [P3, Cor. 4.2].

LEMMA 3.4. *Let T_r be a complete balanced b -ary tree of depth r , and let each vertex of T_r (except for the root) be considered idle with probability q independently. Let the random variable Z_r denote the number of leaves (vertices at depth r) for which every vertex on the path from the root (exclusive) to the leaf (inclusive) is idle. Then we have*

$$\Pr[Z_r = 0] = \xi + O(\eta^r)$$

as $r \rightarrow \infty$ with $b \geq 2$ and $1/b < q < 1$ fixed, where ξ is the unique solution of the equation $(1 - q(1 - \xi))^b = \xi$ in the range $0 < \xi < 1$, and $\eta = b(1 - q(1 - \xi))^{b-1} < 1$.

Now set $r = k/2$ and $s = l - k/2$. The paths from an input v to links in rank r of $G_{b,k,l}$ form a tree isomorphic to T_r (if we ignore the directions of the edges), and the paths from links in rank s to an output w form a disjoint tree isomorphic to T_r . Thus the number of links u in rank r for which all the links on the path from v to u are idle is a random variable U with the same distribution as Z_r . Dually, the number of links u in rank s for which all the links on the path from u to w are idle is an independent random variable U' with the same distribution as Z_r . If v and w are linked, then we

must have $U \geq 1$ and $U' \geq 1$, so by Lemma 3.1 we have

$$Q \leq \Pr[U \geq 1, U' \geq 1] = (1 - \xi)^2 + O(\eta^r).$$

This completes the upper bound for Theorem 1.2.

We now turn to the lower bound for Theorem 1.2. Since this result has been proved for $c \leq 2$ in Pippenger [P3], we shall assume that $c > 2$. (This assumption could of course be avoided, but it would require a more complicated choice of parameters and consideration of cases.) For $0 < q < q_c$, there is nothing to prove, since Q is certainly nonnegative. For $q_c < q < 1$, we shall use the following lemma from Pippenger [P3, Lem. 8.1].

LEMMA 3.5. *With Z_r as in Lemma 3.4 and $1 \leq H \leq (bq)^r$, we have*

$$\Pr[Z_r \leq H] \leq \xi + O\left(\left(H/(bq)^r\right)^\alpha\right)$$

as $r \rightarrow \infty$ with $b \geq 2$ and $1/b < q < 1$ fixed, where $\alpha = \log(1/\eta)/\log(bq)$ and η is as in Lemma 3.4.

Supposing that $q_c < q < 1$, we shall define

$$q_* = q_{c-1} q^{1/(c-1)^2}.$$

We observe that $q < 1$ implies $q_* < q_{c-1}$ and that $q_c < q$ implies $q_* < q$.

LEMMA 3.6. *Let $k \rightarrow \infty$ and $l = ck$, with $b \geq 2$, $q_c < q < 1$, and $c > 2$ all fixed. Then for all sufficiently large k , we have*

$$\psi_h(q_*) \leq k$$

for all $0 \leq h \leq l - k$.

Proof. From Proposition 2.4 we have

$$\varphi_h(q_*) = \left(\frac{b-1}{bq_*-1}\right)^2 b^{h-k} q_*^{h+1} + 1 + O(hb^{h-2k} q_*^h) + O(hq_*^k) + O(hq_*^h).$$

Since $q_* < q_{c-1}$ and $h \leq l - k$, each term is $O(1)$, and thus at most k for all sufficiently large k . \square

Let

$$H = \lceil (bq_*)^r \rceil.$$

We observe that v and w will be linked if the following three events occur:

- I. The input v is joined by paths containing only idle links to all the links in a set V containing at least H idle links in rank r .
- II. All the links in a set W containing at least H idle links in rank s are joined by paths containing only idle links to the output w .
- III. There is at least one path containing only idle links from some link in V to some link in W .

By Lemma 3.5, we have

$$\Pr[I] \geq 1 - \xi + O\left(\left(q_*/q\right)^r\right),$$

and since $q_* < q$ we have $\Pr[I] \rightarrow 1 - \xi$. Dually, we have by Lemma 3.5

$$\Pr[II] \geq 1 - \xi + O\left(\left(q_*/q\right)^r\right),$$

and thus also $\Pr[II] \rightarrow 1 - \xi$. Since events I and II are independent, we have $\Pr[I, II] \rightarrow (1 - \xi)^2$. Thus to complete the proof of the lower bound for Theorem 1.2, it will suffice to show that

$$\Pr[III | I, II] \rightarrow 1.$$

Event III depends on events I and II through the sets V and W . We can avoid having to consider this dependence by showing that $\Pr[III] \rightarrow 1$ for *any* sets V and W each containing at least H links. Thus it will suffice to prove the following proposition.

PROPOSITION 3.7. *Let V and W be any sets of links in ranks r and s , respectively, each containing at least H links. Then*

$$\Pr[III] \rightarrow 1$$

as $k \rightarrow \infty$ with $l = ck$, and with $b \geq 2$, $c > 2$, and $q_c < q < 1$ all fixed.

Proof. Since $\Pr[III]$ can only increase if links are added to V or W , we may assume that V and W each contain *exactly* H links. Also, since $\Pr[III]$ can only increase if q is increased, it will suffice to estimate $\Pr[III]$, assuming the vacancy probability to be $q_* < q$ rather than q .

Let the random variable Y be the number of paths containing only idle links joining some link in V (exclusive) to some link in W (exclusive). Then event III is equivalent to $Y \geq 1$ and thus it will suffice to show that $\Pr[Y = 0] \rightarrow 0$. To do this, we shall use Chebyshev's inequality:

$$\Pr[Y = 0] \leq \frac{\text{Var}[Y]}{\text{Ex}[Y]^2}.$$

Each path from a link in rank r (exclusive) to a link in rank s (exclusive) contains $s - r - 1 = l - k - 1$ links. Since each of these links is independently idle with probability q_* , the probability that such a path contains only idle links is q_*^{l-k-1} . By Corollary 3.3, the number of such paths joining a given link in rank r with a given link in rank s is $b^{s-r-k} = b^{l-2k}$. Since there are H links in each of V and W , we have

$$\text{Ex}[Y] = H^2 b^{l-2k} q_*^{l-k-1}.$$

Next we must estimate $\text{Var}[Y]$. We have

$$\begin{aligned} \text{Var}[Y] &= \sum_{u':V \rightarrow W} \sum_{u:V \rightarrow W} (\Pr[u, u' \text{ idle}] - \Pr[u \text{ idle}] \Pr[u' \text{ idle}]) \\ &= \sum_{u':V \rightarrow W} \Pr[u' \text{ idle}] \sum_{u:V \rightarrow W} (\Pr[u \text{ idle} | u' \text{ idle}] - \Pr[u \text{ idle}]). \end{aligned}$$

Here each sum is over all H^2 paths joining a link in V to a link in W , so there are H^4 terms in all. If u is a path from a link in rank r to a link in rank s , let $\rho(u)$ denote the link in rank r and $\sigma(u)$ the link in rank s . By Lemma 2.1, we may assume (as in the proof of Theorem 1.1) that $u' = u^*$ is part of a path from $v^* = 0^k$ through $\rho(u') = 0^k$ and $\sigma(u') = 0^k$ to $w^* = 0^k$, in which all the links have label 0^k . Thus we have

$$\text{Var}[Y] = H^2 b^{l-2k} q_*^{l-k-1} \sum_{u:V \rightarrow W} (\Pr[u \text{ idle} | u^* \text{ idle}] - \Pr[u \text{ idle}]).$$

The factor $H^2 b^{l-2k} q_*^{l-k-1}$ multiplying the sum is $\text{Ex}[Y]$, so to show that

$$\text{Var}[Y]/\text{Ex}[Y]^2 \rightarrow 0,$$

it will suffice to show that $J/\text{Ex}[Y] \rightarrow 0$, where

$$J = \sum_{u:V \rightarrow W} (\Pr[u \text{ idle} \mid u^* \text{ idle}] - \Pr[u \text{ idle}]).$$

We now partition the paths u into four classes as follows:

- i. those for which $\varrho(u) = \sigma(u) = 0^k$;
- ii. those for which $\varrho(u) \neq 0^k$ but $\sigma(u) = 0^k$;
- iii. those for which $\varrho(u) = 0^k$ but $\sigma(u) \neq 0^k$; and
- iv. those for which $\varrho(u) \neq 0^k$ and $\sigma(u) \neq 0^k$.

We shall denote the contributions to J over these four classes by J_i , J_{ii} , J_{iii} , and J_{iv} , respectively, and estimate them in turn.

For J_i , we have

$$\begin{aligned} J_i &\leq \sum_{u:0^k \rightarrow 0^k} \Pr[u \text{ idle} \mid u^* \text{ idle}] \\ &= \varphi_{s-r}(q_*) \\ &\leq k \end{aligned}$$

by Lemma 3.6. Thus we have

$$\begin{aligned} \frac{J_i}{\text{Ex}[Y]} &\leq \frac{k}{H^2 b^{l-2k} q_*^{l-k-1}} \\ &\leq \frac{k}{b^{l-k} q_*^l} \\ &\rightarrow 0, \end{aligned}$$

since $q_* > q_c$.

For J_{ii} , we have

$$J_{ii} \leq \sum_{V \setminus \{0^k\} \rightarrow 0^k} \Pr[u \text{ idle} \mid u^* \text{ idle}].$$

To estimate $\Pr[u \text{ idle} \mid u^* \text{ idle}]$, let i be the first rank for which a link in u has label 0^k . Since there are two distinct paths in $\langle G_{b,k,l} \rangle_{0,i}$ from v^* through $\varrho(u^*) = 0^k$ and $\varrho(u) \neq 0^k$ to this link, we must have $i \geq k+1$ by Corollary 3.3. Thus we have

$$\begin{aligned} J_{ii} &\leq (H-1) \left(\sum_{k+1 \leq i \leq k+r} q_*^{i-r-1} \varphi_{s-i}(q_*) + \sum_{k+r+1 \leq i \leq s} b^{i-r-k} q_*^{i-r-1} \varphi_{s-i}(q_*) \right) \\ &\leq (H-1)k \left(\sum_{k+1 \leq i \leq k+r} q_*^{i-r-1} + \sum_{k+r+1 \leq i \leq s} b^{i-r-k} q_*^{i-r-1} \right), \end{aligned}$$

where the factor of $H-1$ accounts for the choice of $\varrho(u) \in V \setminus \{0^k\}$, the factors preceding $\varphi_{s-i}(q_*)$ in the sums account for the probability that all the links on u between ranks r (exclusive) and i (exclusive) are idle, the factors of $\varphi_{s-i}(q_*)$ account for the probability that all the links of u between ranks i and s that are not labeled 0^k are idle, and we have bounded $\varphi_{s-i}(q_*)$ using Lemma 3.6. Bounding the sums by the

number of terms (at most $s - r = l - k$) times the largest term (the first for the first sum, and the last for the second), we have

$$J_{ii} \leq (H - 1)k(l - k)(q_*^{k/2} + b^{l-2k}q_*^{l-k-1}).$$

Thus we have

$$\begin{aligned} \frac{J_{ii}}{\text{Ex}[Y]} &\leq \frac{k(l - k)(q_*^{k/2} + b^{l-2k}q_*^{l-k-1})}{H b^{l-2k}q_*^{l-k-1}} \\ &\leq k(l - k) \left(\frac{1}{b^{l-3k/2}q_*^{l-k}} + \frac{1}{(bq_*)^{k/2}} \right) \\ &\rightarrow 0, \end{aligned}$$

since $b^{c-3/2}q_*^{c-1} > b^{c-3/2}q_c^{c-1} = b^{-1/2}q_c^{-1} > b^{-1/2}q_2^{-1} = 1$ (because $q_* > q_c$, $b^{c-1}q_c^c = 1$, $q_c < q_2$, and $bq_2^2 = 1$) and $bq_* > 1$ (because $q_* > q_\infty = 1/b$).

Dually, we have $J_{iii}/\text{Ex}[Y] \rightarrow 0$.

Finally, for J_{iv} we have

$$\begin{aligned} J_{iv} &= \sum_{u: V \setminus \{0^k\} \rightarrow W \setminus \{0^k\}} (\Pr[u \text{ idle} \mid u^* \text{ idle}] - \Pr[u \text{ idle}]) \\ &= \sum_{\substack{u: V \setminus \{0^k\} \rightarrow W \setminus \{0^k\} \\ u \cap u^* \neq \emptyset}} (\Pr[u \text{ idle} \mid u^* \text{ idle}] - \Pr[u \text{ idle}]) \\ &\leq \sum_{\substack{u: V \setminus \{0^k\} \rightarrow W \setminus \{0^k\} \\ u \cap u^* \neq \emptyset}} \Pr[u \text{ idle} \mid u^* \text{ idle}], \end{aligned}$$

since if $u \cap u^* = \emptyset$, the events “ u idle” and “ u^* idle” are independent, and the summand $\Pr[u \text{ idle} \mid u^* \text{ idle}] - \Pr[u \text{ idle}]$ vanishes. Given a path u with $u \cap u^* \neq \emptyset$, let i be the first rank in which u has a link with label 0^k , and let $j \geq i$ be the last such rank. As in case ii, we have $k + 1 \leq i$, and dually we have $j \leq l - k - 1$. Thus we have

$$\begin{aligned} J_{iv} &\leq (H - 1)^2 \left(\sum_{k+1 \leq i \leq k+r} \sum_{\substack{l-k-r \leq j \leq l-k-1 \\ i \leq j}} q_*^{i-r-1} \varphi_{j-i}(q_*) q_*^{s-j-1} \right. \\ &\quad + \sum_{k+1 \leq i \leq k+r} \sum_{\substack{r \leq j \leq l-k-r-1 \\ i \leq j}} q_*^{i-r-1} \varphi_{j-i}(q_*) b^{s-j-k} q_*^{s-j-1} \\ &\quad + \sum_{k+r+1 \leq i \leq s} \sum_{\substack{l-k-r \leq j \leq l-k-1 \\ i \leq j}} b^{i-r-k} q_*^{i-r-1} \varphi_{j-i}(q_*) q_*^{s-j-1} \\ &\quad \left. + \sum_{k+r+1 \leq i \leq s} \sum_{\substack{r \leq j \leq l-k-r-1 \\ i \leq j}} b^{i-r-k} q_*^{i-r-1} \varphi_{j-i}(q_*) b^{s-j-k} q_*^{s-j-1} \right). \end{aligned}$$

Here we have broken the sum into four parts, according to whether $k + 1 \leq i \leq k + r$ or $k + r + 1 \leq i \leq s$, and also according to whether $l - k - r \leq j \leq l - k - 1$ or $r \leq j \leq l - k - r - 1$. (We note that the second and third double sums will vanish unless

$c > 5/2$, and the fourth double sum will vanish unless $c > 3$.) The factor of $(H - 1)^2$ accounts for the choice of $\varrho(u) \in V \setminus \{0^k\}$ and $\sigma(u) \in W \setminus \{0^k\}$, the factors preceding $\varphi_{j-i}(q_*)$ in the summands account for the probability that the links of u in ranks less than i are idle, the factors of $\varphi_{j-i}(q_*)$ account for the probability that the links of u between i and j and not labeled 0^k are idle, and the factors following $\varphi_{j-i}(q_*)$ in the summands account for the probability that the links of u in ranks greater than j are idle. Bounding the factors $\varphi_{j-i}(q_*)$ using Lemma 3.6, and bounding each double summation by the number of terms (at most $(l - k)^2$) times the largest term (which occurs for $i = k + 1$ and $j = l - k - r$ in the first sum, and for $i = j$ in the remaining three sums), we obtain

$$J_{\text{iv}} \leq (H - 1)^2 k(l - k)^2 \left(q_*^k + 2b^{l-5k/2-1} q_*^{l-k-2} + b^{l-3k} q_*^{l-k-2} \right).$$

Thus we have

$$\begin{aligned} \frac{J_{\text{iv}}}{\text{Ex}[Y]} &\leq k(l - k)^2 \left(\frac{1}{(bq_*)^{l-2k}} + \frac{2}{q_* b^{k/2+1}} + \frac{1}{q_* b^k} \right) \\ &\rightarrow 0, \end{aligned}$$

since $bq_* > 1$, $c > 2$, and $b \geq 2$. This completes the proof of the proposition, and with it the proof of Theorem 1.2. \square

4. Conclusion. We have determined the limiting value of the linking probability in spider-web networks with scale k and depth l when $l = ck$ with $c > 1$. The same method could be used when $l/k \rightarrow \infty$ but $(\log(l + 1))/(k + 1) \rightarrow 0$. In this case, the phase transition would be less abrupt: the limiting value of Q , and even its first derivative with respect to q , would be continuous at the critical value $q_\infty = 1/b$, but the second derivative would be discontinuous. Little would be gained by such networks, however, over those with a large fixed value of c : Their great cost would decrease the critical vacancy probability through only a small interval $[q_\infty, q_c]$, and would provide only a small linking probability in this interval.

Another extension of our results would be to consider, instead of the “independent” probability distribution on states introduced by Lee [L1] and Le Gall [L2, L3], the “coherent” distribution introduced by Pippenger [P1]. (The similar distribution introduced by Koverninskiĭ [K] does not have an obvious generalization for $c > 2$, and in any case it does not seem likely that the additional independence in Koverninskiĭ’s model would have much effect on its tractability for $c > 2$.)

Yet another line of inquiry would be to consider the computational complexity of path-search problems for spider-web networks with $c > 2$, using the link-probe model introduced by Lin and Pippenger [L4]. Such results were obtained by Pippenger [P4] for $c = 2$ (and these results are easily extended to the case $1 < c < 2$), but even for $c = 2$ the known results are far from definitive.

Acknowledgments. The results reported in this paper were obtained during the fourth meeting of the Institute for Elementary Studies, the Focused Research Group on Problems in Discrete Probability, held July 12–26, 2003, at the Banff International Research Station in Banff, Alberta, Canada. The author is especially grateful to Yuval Peres, one of the organizers of that meeting, for urging continued faith in the power of the second-moment method.

REFERENCES

- [B] S. R. BROADBENT AND J. M. HAMMERSLEY, *Percolation processes. I. Crystals and mazes*, Proc. Cambridge Philos. Soc., 53 (1957), pp. 629–641.
- [C] F. R. K. CHUNG AND F. K. HWANG, *The connection patterns of two complete binary trees*, SIAM J. Alg. Disc. Meth., 1 (1980), pp. 322–335.
- [G] G. GRIMMETT, *Percolation*, 2nd ed., Springer-Verlag, Berlin, 1999.
- [I] N. IKENO, *A limit of crosspoint number*, IEEE Trans. Circuit Theory, 6 (1959), pp. 187–196.
- [K] I. V. KOVERNINSKIĬ, *Estimation of the blocking probability for switching circuits by means of probability graphs*, Problems Inform. Transmission, 11 (1975), pp. 63–71.
- [L1] C. Y. LEE, *Analysis of switching networks*, Bell System Tech. J., 34 (1955), pp. 1287–1315.
- [L2] P. LE GALL, *Étude du blocage dans les systèmes de commutation téléphoniques automatiques utilisant des commutateurs électroniques du type crossbar*, Ann. Télécommun., 11 (1956), pp. 159–171; 180–194; 197.
- [L3] P. LE GALL, *Méthode de calcul de l'encombrement dans les systèmes téléphoniques automatiques a marquage*, Ann. Télécommun., 12 (1957), pp. 374–386.
- [L4] G. LIN AND N. PIPPENGER, *Routing algorithms for switching networks with probabilistic traffic*, Networks, 28 (1996), pp. 21–29.
- [P1] N. PIPPENGER, *On crossbar switching networks*, IEEE Trans. Comm., COM-23 (1975), pp. 646–659.
- [P2] N. PIPPENGER, *The blocking probability of spider-web networks*, Random Structures Algorithms, 2 (1991), pp. 121–149.
- [P3] N. PIPPENGER, *The asymptotic optimality of spider-web networks*, Discrete Appl. Math., 37/38 (1992), pp. 437–450.
- [P4] N. PIPPENGER, *Upper and lower bounds for the average-case complexity of path-search*, Networks, 33 (1999), pp. 249–259.
- [T] K. TAKAGI, *Design of multi-stage link systems by means of optimum channel graphs*, Electron. Commun. Japan, 51 (1968), pp. 37–46.