Available online at www.sciencedirect.com
SCIENCE DIRECT®
Theoretical Computer Science 333 (2005) 415–432
Theoretical Computer Science
www.elsevier.com/locate/tcs

Interval routing in reliability networks☆

Cyril Gavoille[a,*], Martin Nehéz[b,1]

[a] *Laboratoire Bordelais de Recherche en Informatique, Université Bordeaux I, 351, cours de la Liberation, Talence Cedex 33405, France*

[b] *Faculty of Informatics and Information Technologies, Slovak University of Technology, Ilkovičova 3, 842 16 Bratislava*

Received 8 September 2003; received in revised form 8 April 2004

Abstract

In this paper, we consider routing with compact tables in reliability networks. More precisely, we study interval routing on random graphs $\mathbb{G}(B, p)$ obtained from a base graph B by independently removing each edge with a failure probability $1-p$. We focus on additive stretched routing for n-node random graphs for which the base B is a square mesh and $p=0.5$, that is the percolation model at the critical phase. We show a lower bound of $\Omega(\sqrt{\log n}/(\delta+2))$ on the number of intervals required per edge for every additive stretch $\delta \geqslant 0$. On the other side, our experimental results show that the size of the largest biconnected components is $\Theta(n^{0.827})$, and thus that there exists a trivial shortest-path routing scheme using at most $O(n^{0.827})$ intervals per edge.

The results are extended to random meshes of higher dimension. We show that, asymptotically almost surely, the number of intervals per edge for a random r-dimensional mesh with n nodes is $\Omega(16^{-r}(\delta+2)^{1-r}r^{-4}(\log n)^{1-1/r})$, for every additive stretch $\delta \geqslant 0$ and for every integral dimension $r \in [1, \log_2 n]$.

Keywords: Compact routing tables; Reliability networks; Random graphs; Interval routing; Percolation theory

☆ The preliminary version of this paper has been presented at the conference SIROCCO 2003.

* Corresponding author.
E-mail addresses: gavoille@labri.fr (C. Gavoille), nehez@fiit.stuba.sk (M. Nehéz).

[1] Supported by the VEGA Grants VG 1/0162/03 and VG 1/0263/03.

0304-3975/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2004.12.008

1. Introduction

Research in the area of routing algorithms on computer networks is permanently of great interest by many researchers since routing problems belong to the fundamental topics in distributed systems. Routing is interesting both from a theoretical and also from a practical point of view. One aspect of research in this field covers compacting of routing tables, by maintaining the smallest amount of routing information (or knowledge) locally in each router while guaranteeing that the routes are near the shortest paths.

Many results concern the design of *universal* routing strategies in the sense that they are applicable to all the networks. In particular, the proposed schemes give trade-offs between the memory requirements (the size of the local routing tables) and the stretch factor, namely the maximum ratio between the length of the route between any two nodes and their distance in the network. Among them [3,5,9,32,38,13] are for a survey.

While the above strategies apply to all the networks, a natural question is whether other more efficient techniques can be applied on realistic networks. Although there is still no answer to the question "what a realistic network is?", many models consider that such networks are based on some structured underlying topology (which is certainly not the complete network) with some random extra connections or some random link failures (cf. the augmented grid Kleingberg's model of small world [23,24]).

1.1. Reliability networks

In this paper, we consider a point-to-point communication network modeled by a simple connected graph $G = (V, E)$, where V is the set of nodes (or processors or routers) and E is a set of edges (or bidirectional communication links). We focus on random graphs $\mathbb{G}(B, p)$ obtained from a graph B, with node set $\{1, \ldots, |V(B)|\}$, by independently removing each edge with a failure probability $1-p$. So $G \in \mathbb{G}(B, p)$ is a uniform labeled random spanning subgraph of B as $V(G) = V(B)$ and $E(G) \subseteq E(B)$. More precisely,

$$\Pr(G) = p^{|E(G)|}(1-p)^{|E(B)|-|E(G)|}.$$

The graph B is called the *base* graph, and the value $1-p$ the *failure probability*. This model, called the *reliability network model*, appears in [28–30] and is described in more detail in [22, p. 2]. The reliability network is a natural generalization of the binomial random graph model of Erdös-Rényi, denoted hereafter $\mathbb{G}(n, p)$, for which $B = K_n$ is the complete graph on n nodes. The reliability network based on the infinite square mesh represents the square bond percolation model described in [21]. As mentioned in [22], this model can be generalized further by allowing different probabilities of failure at different edges. It is also related to other problems of computer science such as grid-computing, fault-tolerant distributed computing, effective data structures, etc.

Our paper is also concerning *additive* stretched routing schemes, rather than multiplicative stretch (or stretch factor). A path of a graph is *δ-stretched* if the length of the path is at most the length of a shortest path between its extremities plus δ. A *δ-stretched routing scheme* is a scheme for which all the routes are δ-stretched paths. The parameter δ is also called the *deviation* of the routing scheme. It is provable that even a small deviation allows better

optimizations for spanner construction [7] and distance computation [16], and yields also compact routing tables [4].

1.2. Routing tables and interval routing

Finally, we focus on the *interval routing scheme*, a particular way of implementing standard *routing tables* [35,40]. Each node has a local routing table, and addresses of the nodes range in $\{1, \ldots, |V(G)|\}$. When a source u sends a message to a destination v, it attaches to the message the address of v, say the integer i, and forwards i and the message to a neighbor of u by looking at the ith entry of u's local table. So the route is computed in a distributed fashion by the nodes along the route between u and v. Obviously, it is required that for every source–destination pair u, v a route connects u to v. Interval routing implements local routing tables as follows: u stores a d entry table, d being the degree of u. Each entry corresponds to the list of destination addresses using the same first edge in the routing from u. If for every node u, all the lists can be grouped into at most k sets of consecutive integers (consecutive modulo n), we say that the routing scheme is a *k-interval routing scheme* (k-IRS for short). For more precise formulations and other details see [12,34].

The main difficulty in the design of interval routing schemes for a given graph is to find out a suitable address assignment for the nodes and a suitable system of routes for all the pairs of nodes such that the number of intervals per output port (equivalently per outgoing edge) is minimal while keeping the routes near-shortest paths. Whenever shortest paths are required, the problem to know whether a graph supports a 1-IRS is already NP-complete [6]. A δ-stretched k-IRS, denoted by (k, δ)-IRS, is simply a k-IRS that is a δ-stretched routing scheme. A $(k, 0)$-IRS is also called a shortest-path k-IRS. Fig. 1 depicts two interval routing schemes on the 6-cycle.

1.3. Previous works on random graphs

The main advantage of k-IRS concerns the size of the memory requirements. In an n-node graph supporting a k-IRS, a node of degree d has to store $O(kd \log n)$ bits of information, whereas $\Omega(n \log d)$ bits are required for a standard routing table implementation.

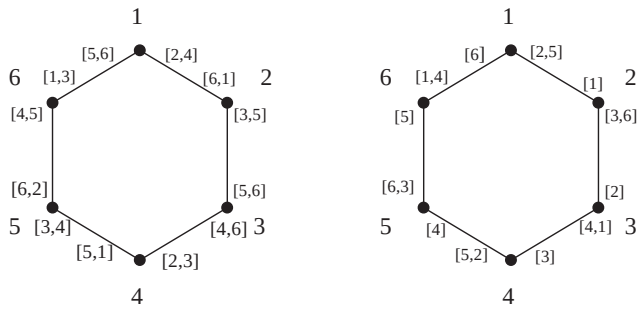

Fig. 1. A shortest-path 1-IRS (left side), and a 2-stretched 1-IRS (right side).

In particular, interval routing is efficient for structured graphs like cycles, complete graphs, meshes, trees, outerplanar graphs, tori, hypercubes, k-trees, etc. All these graphs support shortest-path O(1)-IRS (more results about interval routing are accessible in the survey [12]).

Flammini, van Leeuwen and Marchetti-Spaccamela [8] proved a non-constant lower bound on k for shortest-path k-IRS on random graphs of $\mathbb{G}(n, p)$, the Erdös–Rényi model. It is proved therein that, with high probability, a graph $\mathbb{G}(n, p)$ requires $\Omega(n^{1-1/\Theta(\sqrt{\log n})})$ intervals per outgoing edge for some specific value of p, namely for $p = 1/n^{1-1/\Theta(\sqrt{\log n})}$. On the other hand, Gavoille and Peleg [18] proved that almost all graphs (that is a fraction of $1 - o(1)$ of all n-node graphs, or equivalently the graphs of $\mathbb{G}(n, p)$ for $p = 0.5$ and with high probability) support a shortest-path 2-IRS. Actually, they constructed a routing scheme such that every node has at most $O(\log^3 n)$ outgoing edges with 2 intervals, all the other ones having 1 interval, leaving open the question of whether almost all graphs support shortest-path 1-IRS. Finally, for shortest-path k-IRS on random n-node tori (i.e. Cartesian product of two cycles, each of $\sqrt{n}$ nodes with random deletion of edges), a preliminary result appears in [27] where it is proved a lower bound of $\Omega(\sqrt{\log n})$. A summary of these results is listed in the following table.

Random graph	Probability p	Shortest-path k-IRS	Reference
$G \in \mathbb{G}(n, p)$	$1/n^{1-1/\Theta(\sqrt{\log n})}$	$\Omega\left(n^{1-1/\Theta(\sqrt{\log n})}\right)$	[8]
$G \in \mathbb{G}(n, p)$	constant	$\leqslant 2$	[18]
$T \in \mathbb{G}\left(T_{\sqrt{n}\times\sqrt{n}}, p\right)$	constant	$\Omega\left(\sqrt{\log n}\right)$	[27]

1.4. Our results

The main results of this paper are the following:

1. A lower bound on k for (k, δ)-IRS on random n-node r-dimensional meshes with constant failure probability. We show that, asymptotically almost surely, $k = \Omega(16^{-r}(\delta + 2)^{1-r} r^{-4} (\log n)^{1-1/r})$, for every additive stretch $\delta \geqslant 0$ and for every dimension $r \in [1, \log_2 n]$.
2. For upper bounds, we have studied random square meshes ($r = 2$) with $p = 0.5$, that is the percolation threshold probability. Recall that when the size of the mesh becomes infinite, $p = 0.5$ is precisely the probability where the mesh contains a unique infinite connected component [21]. Unfortunately, as many interesting problems in percolation theory (and as suggested by [1]), few answers can be done analytically and we restrict ourselves to make experiments. Based on the expected size of the largest biconnected components of random square meshes, our results suggest that random square n-node meshes support shortest-path k-IRS with $k = O(n^{0.827})$.

The motivation for studying, meshes is that meshes or the subgraphs of a mesh are typical planar graphs (cf. the graph minor theory of Robertson and Seymour [33]) and many problems are still unsolved about routing in planar graphs with compact tables. For instance, the optimal size of shortest-path routing tables is not known. The complexity bound

ranges between $\Omega(\sqrt{n})$ [5] and $O(n)$ bits per node [15,26]. For shortest-path interval routing the range is similar: $\Omega(\sqrt{n})$ intervals is the best known lower bound [19,39] and $O(n)$ is the trivial upper bound. Similar gaps exist also for *distance labeling* in planar graphs whose goal is to compute distances between two nodes based only on their node label [31]: label length must be $\Omega(n^{1/3})$ for some worst-case, and $O(\sqrt{n}\log n)$ bit labels are sufficient for every planar graph [20]. Finding structure of shortest paths and distances in planar graphs is probably difficult and certainly would require more combinatorics.

Surprisingly, bounds for multiplicative stretched routing and approximated distance labeling are much more competitive. In general almost poly-logarithmic space per node is sufficient [11,37]. Observe that shortest-path and additive stretched routing in planar graphs are two equivalent problems in the sense that the lower bounds on the shortest-path version transfer to lower bounds on stretched version by subdividing each edge into δ other edges. This forces any additive $O(\delta)$-stretched routing scheme to respect the shortest paths in the subdivided graph, and this latter graph remains of linear size for constant δ.

The paper is organized as follows: Section 2 presents the lower bound, and Section 3 the upper bound and our of experiments. We conclude by a large set of open problems in Section 4.

2. Lower bound for random r-dimensional meshes

If G is a connected graph, then we denote by $\mathrm{IRS}_\delta(G)$ the smallest integer k such that G supports a (k,δ)-IRS. The number $\mathrm{IRS}_0(G)$ is also called the *compactness* of G. Because the graphs of $\mathbb{G}(B,p)$ are not necessarily connected, we extend the notion of the routing schemes on non-connected graphs as follows: a routing scheme on a non-connected graph is simply the union of the routing schemes of each of its connected components. It is therefore only required to have a route between two nodes of a same connected component. Then, $\mathrm{IRS}_\delta(G) = \max_i \mathrm{IRS}_\delta(G_i)$ where the G_is are the connected components of G.

For two integers $r \geqslant 1$ and $s \geqslant 1$, the *r-dimensional mesh*, denoted M_s^r, is the graph whose nodes are all the r-tuples over the set $\{1,\ldots,s\}$. Two nodes $u = (u_1,\ldots,u_r)$ and $v = (v_1,\ldots,v_r)$ are adjacent if and only if there is an index i_0 such that $|u_{i_0} - v_{i_0}| = 1$ and $u_i = v_i$ for every $i \neq i_0$. The graph M_s^r has s^r nodes and $rs^{r-1}(s-1)$ edges.

This section concerns graphs of $\mathbb{G}(M_s^r, p)$. In order to prove the lower bound we need several preliminary results.

A subgraph X of G is *isolated* if there is no edge $\{u,v\} \in E(G)$ such that $u \in V(X)$ and $v \notin V(X)$. Note that if X is isolated in G, then $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(X)$ for every $\delta \geqslant 0$, since by definition $\mathrm{IRS}_\delta(G) = \max_i \mathrm{IRS}_\delta(G_i)$.

A graph X is an *m-subgraph of G* if G contains m subgraphs isomorphic to X pairwise at distance two or more. For instance, K_2 is a 2-subgraph of the 6-cycle depicted in Fig. 1 (take two opposite edges). We emphasize that if X is an m-subgraph of G then X is not necessarily an induced subgraph of G.

Lemma 1. *Let B be a connected graph with n nodes and maximum degree d, and let X be an m-subgraph of B with x nodes. Then, $G \in \mathbb{G}(B,p)$ contains X as isolated subgraph with probability at least $1 - \exp(-mq^{dx})$, where $q = \min\{p, 1-p\}$.*

Proof. Let $G \in \mathbb{G}(B, p)$, and let A be the event "G contains X as isolated subgraph". Our goal is to lower bound $\Pr(A)$. So,

$$\Pr(A) \geqslant \Pr(X \text{ is isolated in } G).$$

Let $X_1, \ldots, X_m$ be m subgraphs of B pairwise at distance two, each one isomorphic to X. For each $i \in \{1, \ldots, m\}$, let Z_i be the random boolean variable such that $Z_i = 1$ if and only if the subgraph X_i is isolated in G. Finally, let $Z = \sum_{i=1}^{m} Z_i$.

$$\Pr(X \text{ is isolated in } G) \geqslant \Pr(Z \geqslant 1) = 1 - \Pr(Z = 0).$$

As the X_i's are pairwise at distance at least two, there is no edge $\{u, v\} \in E(B)$ with $u \in X_i$ and $v \in X_j$. Therefore, "X_i is isolated in G" is an event independent from "X_j is isolated in G". Thus the variables Z_i are mutually independent. It follows that,

$$\Pr(Z = 0) = \prod_{i=1}^{m} \Pr(Z_i = 0) \leqslant \left(\max_i \Pr(Z_i = 0)\right)^m = \left(1 - \min_i \Pr(Z_i = 1)\right)^m.$$

To make X_i isolated in B it suffices to keep or remove independently some edges of B with at least one extremity is in X_i. It is possible to check (recall that the degree of each node of B is bounded by d) that $\Pr(Z_i = 1) \geqslant q^{dx}$, where $q = \min\{p, 1 - p\}$. Thus,

$$\Pr(A) \geqslant \Pr(Z \geqslant 1) \geqslant 1 - \left(1 - q^{dx}\right)^m \geqslant 1 - \exp\left(-mq^{dx}\right)$$

using the fact that $(1 - b/a)^c \leqslant \exp(-c(b/a))$, for all $0 < b \leqslant b + c < a$ (cf. [2, Eq. (1.6) p. 5]). This completes the proof. □

In the following, we denote by Π a *routing property*, that is a set of possible routes for a routing scheme on a graph. More formally, a routing property Π is a function that associates with every graph G a set $\Pi(G)$ of paths of G. A routing scheme R on G has the property Π (or is a Π-routing) if all the routes induced by R belong to $\Pi(G)$. For instance, the "shortest-path" property is simply a function Π such that, for every G, $\Pi(G)$ returns the set of all the shortest paths in G.

The following useful lemma is a generalization of a result of Kráľovič et al. [25] originally proved for the shortest-path property.

Lemma 2. *Let G be a graph and Π be a routing property. Let $S(u, v)$ be the set of nodes w such that there exists a path of $\Pi(G)$ from u to w that starts with the edge $\{u, v\}$. Let U and W be two disjoint node subsets of G such that for all distinct nodes $w, w' \in W$, there is $u \in U$ such that for each neighbor v of u it holds $w \notin S(u, v)$ or $w' \notin S(u, v)$. Then, every k-IRS with property Π on G must satisfy*

$$k \geqslant \frac{|W|}{\sum_{u \in U} \deg(u)}.$$

Proof. Let R be any k-IRS with property Π on a graph G with node set $\{v_1, \ldots, v_n\}$ and arc set (each edge appears twice, once for each orientation) $\{e_1, \ldots, e_m\}$. From R we construct a $n \times m$ boolean matrix $M_{i,j}$ as follows: $M_{i,j} = 1$ iff the route induced by R from u_j (the

tail of e_j) to v_i starts with the edge e_j, and set $M_{i,j} = 0$ otherwise. It is not difficult to see that the number of intervals associated with the edge e_j by R is exactly the number of 01-sequences in the binary vector composed of jth column of $M_{i,j}$ (the last bit and the first bit being considered as consecutive). Let $c(e_j)$ be its 01-sequences number. By the choice of R, $c(e_j) \leqslant k$ for every j.

Consider now the sub-matrix M' composed of all the rows corresponding to a node of W and of all the columns corresponding to an arc outgoing from a node of U. Let $c'(e_j)$ be the 01-sequences number of the column e_j of M'. Removing some bits of a binary vector does not increase its 01-sequence number. Hence, $c'(e_j) \leqslant c(e_j)$ and so $k \geqslant c'(e_j)$.

Let us show that $\sum_j c'(e_j) \geqslant |W|$, where the sum is done over all the columns of M'. Indeed, consider two consecutive rows of M' (again the last and the first row are considered as consecutive), and let w, w' be the corresponding nodes of W, say w' located below w in M'. Consider the node $u \in U$ such that for each neighbor v of u, $w \notin S(u, v)$ or $w' \notin S(u, v)$. R has the property Π, thus u has a neighbor v' such that $w' \in S(u, v')$. As $w \notin S(u, v')$, it follows that the binary vector of the column associated with the arc (u, v') contains a 01-sequence starting at the row of w. Considering all the consecutive pairs w, w' of M' we have that all the columns of M' contain at least $|W|$ disjoint 01-sequences, i.e., $\sum_j c'(e_j) \geqslant |W|$ as claimed.

As the number of columns of M' is $\sum_{u \in U} \deg(u)$, it follows that

$$k \geqslant \frac{\sum_j c'(e_j)}{\sum_{u \in U} \deg(u)} \geqslant \frac{|W|}{\sum_{u \in U} \deg(u)}. \qquad \square$$

In Lemma 3, we will apply Lemma 2 where $\Pi(G)$ is the set of δ-stretched paths on G.

For a positive integer t, let us define an operation $\star$ over a graph G and the t-node path P_t as follows: $G \star P_t$ is a graph which consists of two copies of G, namely G and G', such that each node $u \in G$ is connected with its corresponding copy $u' \in G'$ by a path P_t. For all $t, r \geqslant 1$, we define the graph H_t^r recursively by $H_t^{r+1} = H_t^r \star P_t$, with $H_t^1 = P_t$. In other words,

$$H_t^r = \underbrace{P_t \star P_t \star \cdots \star P_t}_{r \text{ times}}.$$

Note that H_t^r is a subgraph of the r-dimensional mesh M_s^r for $t \leqslant s$. Similarly as in meshes M_s^r, we will identify nodes of H_t^r by r-tuples over the set $\{1, \ldots, t\}$. The number of nodes of H_t^r is t^r, and its number of edges is $(t-1)(t^r - 2^r)/(t-2)$, that follows from the recurrence: $|E(H_t^{i+1})| = 2|E(H_t^i)| + (t-1)t^i$ with $|E(H_t^0)| = 0$. See Fig. 2 for some examples of H_t^r.

Lemma 3. $\mathrm{IRS}_\delta(H_t^r) = \Omega\left((\delta+2)^{1-r} r^{-2} 2^{-r} (t-7)^{r-1}\right)$, *for all* $\delta \geqslant 0$, $r \geqslant 1$ *and* $t \geqslant 8$.

Proof. The result clearly holds if $\delta \geqslant t$, since in this case $\mathrm{IRS}_\delta(H_t^r) = \Omega\left(r^{-2} 2^{-r}\right)$. So assume $\delta \leqslant t-1$.

Let U and W be the following sets.

$u \in U$ if in $u = (u_1, \ldots, u_r)$ there is at most one $u_i \neq 1$ ($i \in \{1, \ldots, r\}$).

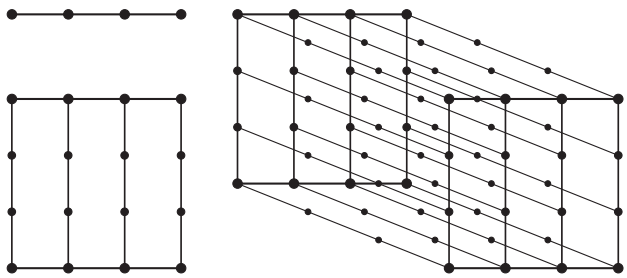

Fig. 2. H_4^1, H_4^2 and H_4^3.

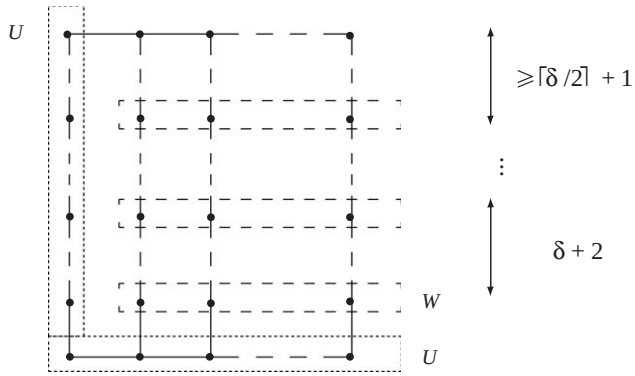

Fig. 3. The sets W and U in the graph H_t^2.

$w \in W$ if $w = (w_1, \ldots, w_r)$, where:

- $2 \leqslant w_1 \leqslant t$,
- $w_i = 2 + l(\delta + 2)$ for $l = 0, \ldots, \lambda - 2$, such that $i \in \{2, \ldots, r\}$ and

$$\lambda = \left\lfloor \frac{t - \lceil \delta/2 \rceil - 3}{\delta + 2} \right\rfloor + 2.$$

The construction of W yields the important fact that all nodes from W are of degree two and have only edges in the rth dimension. Note that for each w_i, $i \in \{2, \ldots, r\}$ it holds that $2 \leqslant w_i \leqslant t - \lceil \delta/2 \rceil - 1$ and $\sum_{u \in U} \deg(u) = O(r^2 t)$, since $|U| \leqslant rt$ and $|W| = (t-1) \cdot (\lambda - 1)^{r-1}$. Clearly, U and W are disjoint and $|W| > 0$. (For the construction of W and U in H_t^2 see Fig. 3.)

Let $\Pi(H_t^r)$ be the δ-stretched path routing property. We will show that the sets U and W satisfy the assumption of Lemma 2. The main idea comes from the following fact: for all pairs of nodes $w, w' \in W$ a node $u \in U$ exists such that there is only one δ-stretched path $P(w, u)$ from w to u and also only one δ-stretched path $P(w', u)$ from w' to u. Moreover paths $P(w, u)$ and $P(w', u)$ are disjoint except for the node u. It follows that there is only one way to reach nodes w and w' from u according to $\Pi(H_t^r)$, and the two corresponding paths start with different edges incident to u, thus satisfying the assumption of Lemma 2.

Let $w = (w_1, \ldots, w_r)$, $w' = (w'_1, \ldots, w'_r)$ be two distinct nodes from W. We will distinguish two cases for setting the node u.

Case 1: Nodes w and w' differ in the first element, i.e. $w = (w_1, \ldots, w_r)$, $w' = (w'_1, \ldots, w'_r)$ and $w_1 \neq w'_1$. Let us choose the corresponding node from U to be $u = (w_1, 1, \ldots, 1)$. Recall that all nodes $w, w' \in W$ are incident only with edges in the rth dimension. Let $P(w, u)$ be the path from w to u containing the following nodes: $(w_1, \ldots, w_{r-1}, 1)$, $(w_1, \ldots, w_{r-2}, 1, 1)$, $\ldots$, $(w_1, 1, \ldots, 1)$. The length $\ell(P(w, u))$ of $P(w, u)$ is given by

$$\ell(P(w, u)) = \sum_{i=2}^{r} (w_i - 1).$$

Hence $P(w, u)$ is the shortest path, since the distance of w and u in H_t^r is equal to $\ell(P(w, u))$. Clearly, $P(w, u)$ satisfies the property $\Pi(H_t^r)$. Another possibility for reaching node u from w is to pass through a node $(w_1, \ldots, w_{j-1}, t, w_{j+1}, \ldots, w_r)$ for an aribtrary j such that $2 \leqslant j \leqslant r$. However, each such path does not satisfy the property $\Pi(H_t^r)$. The argument is that the construction of W yields that the distance between any $w = (w_1, \ldots, w_r) \in W$ and the node $(w_1, \ldots, w_{j-1}, t, w_{j+1}, \ldots, w_r)$ is at least $\lceil \delta/2 \rceil + 1$. So the length of any path that contains the node $(w_1, \ldots, w_{j-1}, t, w_{j+1}, \ldots, w_r)$ is raised above $2(\lceil \delta/2 \rceil + 1)$. More precisely, let $P'(w, u)$ denote a path from w to u that contains the node $(w_1, \ldots, w_{r-1}, t)$. (Note that it holds $j = r$ in this case.) For its length $\ell(P'(w, u))$ it holds

$$\ell(P'(w, u)) = \sum_{i=2}^{r-1} (w_i - 1) + t - 1 + t - w_r = \sum_{i=2}^{r} (w_i - 1) + 2t - 2w_r.$$

Let us express the difference $\ell(P'(w, u)) - \ell(P(w, u))$

$$\ell(P'(w, u)) - \ell(P(w, u)) = 2(t - w_r) \geqslant \delta + 2,$$

since $t - w_r \geqslant \lceil \delta/2 \rceil + 1$. Thus $P(w, u)$ is the only one possible path that satisfies $\Pi(H_t^r)$, since the length of any other path from w to u is greater than or equal to $\ell(P'(w, u))$.

An analogous situation holds also for the node $w' = (w'_1, \ldots, w'_r)$: the only routing path $P(w', u) \in \Pi(H_t^r)$ contains the nodes $(w'_1, \ldots, w'_{r-1}, 1)$, $(w'_1, \ldots, w'_{r-2}, 1, 1)$, $\ldots$, $(w'_1, 1, \ldots, 1)$. Assume w.l.o.g. that $w'_1 > w_1$. Then the length $\ell(P(w', u))$ of $P(w', u)$ is equal to the distance between w' and u

$$\ell(P(w', u)) = \sum_{i=2}^{r} (w'_i - 1) + w'_1 - w_1.$$

On the other hand, let $P'(w', u)$ denote the path from w' to u that contains the node $(w'_1, \ldots, w'_{r-1}, t)$. The length $\ell(P'(w', u))$ is expressed as follows:

$$\begin{aligned}\ell(P'(w', u)) &= \sum_{i=2}^{r-1} (w'_i - 1) + t - 1 + t - w'_r + w'_1 - w_1 \\ &= \sum_{i=2}^{r} (w'_i - 1) + 2t - 2w'_r + w'_1 - w_1,\end{aligned}$$

and consequently,

$$\ell(P'(w',u)) - \ell(P(w',u)) = 2(t - w'_r) \geqslant \delta + 2,$$

since $t - w'_r \geqslant \lceil \delta/2 \rceil + 1$.

Now, it is easy to see that $v = (w_1, 2, 1, \ldots, 1)$ is the last node that is contained in the path $P(w,u) \in \Pi(H_t^r)$ before its end-node $u = (w_1, 1, \ldots, 1)$ and $v' = (w_1{+}1, 1, \ldots, 1)$ is the last node that is contained in the path $P(w',u) \in \Pi(H_t^r)$ before its end-node. Consequently, for neighbors v and v' of u it holds: $v \in P(u,w)$, $v' \in P(u,w')$ and $v \neq v'$. Both paths $P(u,w)$ and $P(u,w')$ are the only possible ones satisfying $\Pi(H_t^r)$. Hence, the assumption of Lemma 2 holds.

Case 2: Let $j > 1$ be the index of the first element in which w and w' are different, i.e. $w = (w_1, \ldots, w_{j-1}, w_j, \ldots, w_r)$ and $w' = (w_1, \ldots, w_{j-1}, w'_j, \ldots, w'_r)$. Let us choose the corresponding node from U to be

$$u = \left(\underbrace{1, \ldots, 1}_{j-1}, t + 1 - \left\lfloor \frac{w_j + w'_j}{2} \right\rfloor, \underbrace{1, \ldots, 1}_{r-j} \right).$$

Assume w.l.o.g. that $w_j > w'_j$. By the argument that all nodes of W are incident only with edges in the rth dimension, there exists the path $P(w,u)$ which contains the following nodes: $(w_1, \ldots, w_j, \ldots, w_{r-1}, 1), \ldots, (w_1, \ldots, w_j, 1, \ldots, 1), \ldots, (w_1, \ldots, w_{j-1}, t, 1, \ldots, 1), (w_1, \ldots, 1, t, 1, \ldots, 1), \ldots, (1, \ldots, 1, t, 1, \ldots, 1)$. The fact $(w_1, \ldots, w_{j-1}, t, 1, \ldots, 1) \in P(w,u)$ yields

$$\begin{aligned} \ell(P(w,u)) &= \sum_{i=1}^{j-1} (w_i - 1) + \sum_{i=j+1}^{r} (w_i - 1) + (t - w_j) \\ &\quad + t - \left(t + 1 - \left\lfloor \frac{w_j + w'_j}{2} \right\rfloor \right) \\ &= \sum_{i=1}^{r} (w_i - 1) + t - 2w_j + \left\lfloor \frac{w_j + w'_j}{2} \right\rfloor. \end{aligned}$$

On the other hand, let $P'(w,u)$ denote the path from w to u that contains the node $(w_1, \ldots, w_{j-1}, 1, 1, \ldots, 1)$. It implies

$$\ell(P'(w,u)) = \sum_{i=1}^{r} (w_i - 1) + t + 1 - \left\lfloor \frac{w_j + w'_j}{2} \right\rfloor - 1.$$

The difference $\ell(P'(w,u)) - \ell(P(w,u))$ is

$$\ell(P'(w,u)) - \ell(P(w,u)) = w_j - w'_j \geqslant \delta + 2,$$

that follows from the definition of W and from the fact that $w_j > w'_j$. Similarly as in the previous case, $P(w,u)$ is the only possible path satisfying $\Pi(H_t^r)$.

An analogous argument yields that there is only one possible path $P(w', u) \in \Pi(H_t^r)$ such that $(w_1, \ldots, w_{j-1}, 1, 1, \ldots, 1) \in P(w', u)$. Note that in this case

$$\ell(P(w', u)) = \sum_{i=1}^{r} (w'_i - 1) + t - \left\lfloor \frac{w_j + w'_j}{2} \right\rfloor,$$

since $P(w', u)$ is the shortest path. Let us denote the expression $t + 1 - \lfloor (w_j + w'_j)/2 \rfloor$ by u_j. The construction of the path $P(w, u)$ yields that $v = (1, \ldots, 1, u_j + 1, 1, \ldots, 1)$ is the last node of $P(w, u) \in \Pi(H_t^r)$ before its end-node $u = (1, \ldots, 1, u_j, 1, \ldots, 1)$ and $v' = (1, \ldots, 1, u_j - 1, 1, \ldots, 1)$ is the last node that is contained in the path $P(w', u) \in \Pi(H_t^r)$ before its end-node u. Now it is easy to see that for each neighbor v of u it holds $w \notin S(u, v)$ or $w' \notin S(u, v)$. (The meaning of $S(u, v)$ is the same as in Lemma 2.)

Lemma 2 leads to the following inequality:

$$\mathrm{IRS}_\delta(H_t^r) \geqslant r^{-2} t^{-1} (t-1)(\lambda - 1)^{r-1}.$$

By the choice of λ,

$$\lambda - 1 \geqslant \frac{t - \delta/2 - 4}{\delta + 2} \geqslant \frac{t - 7}{2(\delta + 2)},$$

since $\delta \leqslant t-1$. Moreover, $t^{-1}(t-1) \geqslant 7/8$, since $t \geqslant 8$ and we obtain the resulting asymptotic formula:

$$\mathrm{IRS}_\delta(H_t^r) = \Omega\left(r^{-2} 2^{-r} \left(\frac{t-7}{\delta+2} \right)^{r-1} \right). \qquad \square$$

The main result of this section is the following.

Theorem 1. *Let $0 < p < 1$ be a constant, and let $s \geqslant 2$, $r \geqslant 1$, $\delta \geqslant 0$ be integers. With probability at least $1 - \exp(-(s/2)^{r-1/2})$ a random mesh $G \in \mathbb{G}(M_s^r, p)$ verifies $\mathrm{IRS}_\delta(G) = \Omega(16^{-r}(\delta+2)^{1-r} r^{-3} (\log s)^{1-1/r})$.*

Proof. The result holds for $r = 1$, as $16^{-r}(\delta+2)^{1-r} r^{-3} (\log s)^{1-1/r} = \mathrm{O}(1)$ and since $\mathrm{IRS}_\delta(G)$ is always at least 1. (Notice that every $G \in \mathbb{G}(M_s^1, p)$ is a forest thus satifying $\mathrm{IRS}_0(G) = 1$ by [35].) We note also that for $s = \mathrm{O}(1)$, $16^{-r}(\delta+2)^{1-r} r^{-3} (\log s)^{1-1/r} = \mathrm{O}(1)$, thus the result holds as well. So let us assume that $r \geqslant 2$ and that s is large enough (i.e., s is greater than some fixed constant s_0).

Let us show that, for every $t \leqslant s - 2$, H_t^r is an m-subgraph of M_s^r for $m = \lfloor s/(t+2) \rfloor^r$. We split each path P_s of the construction of M_s^r into $\lfloor s/(t+2) \rfloor$ segments of $t + 2$ nodes. It splits M_s^r into $\lfloor s/(t+2) \rfloor^r$ sub-meshes isomorphic to M_{t+2}^r. Each sub-mesh contains H_t^r as a subgraph (not necessarily induced) that one can place so that any two copies of H_t^r into M_s^r are at distance at least two. By this way we have shown that M_s^r contains $m = \lfloor s/(t+2) \rfloor^r$ subgraphs isomorphic to H_t^r and pairwise at distance two. For $t \leqslant s - 2$, we have $m \geqslant 1$ and H_t^r is an m-subgraph of M_s^r as claimed.

Let us fix $t = (\alpha r \log_2(s/2))^{1/r}$ with $\alpha = -1/(8r^2 \log_2 q)$ for $q = \min\{p, 1-p\}$.[2] Note that $0 < \alpha < 1$. Observe also that, for every $r \geqslant 1$, $r^{1/r} \in [1, e^{1/e}]$. So, $(\alpha r)^{1/r} \leqslant e^{1/e}$, and thus $t \leqslant e^{1/e}(\log_2(s/2))^{1/r} \leqslant s-2$ for s large enough. Let us express $(t-7)^{r-1}$. Note that from Lemma 3 we have the condition $t \geqslant 8$ and consequently, $t - 7 \geqslant t/8$. It follows that $t - 7 \geqslant 8^{-1}(\alpha r \log_2(s/2))^{1/r}$. Now,

$$\begin{aligned}(t-7)^{r-1} &\geqslant 8^{1-r} \cdot (\alpha r) \cdot (\alpha r)^{-1/r} \cdot (\log_2(s/2))^{1-1/r} \\ &\geqslant 8^{1-r} \alpha r \cdot e^{-1/e} \cdot (\log_2(s/2))^{1-1/r}, \qquad (1) \\ &\geqslant -\frac{e^{-1/e}}{8 \log_2 q} \cdot 8^{1-r} r^{-1} (\log_2(s/2))^{1-1/r} \\ &\geqslant -\frac{e^{-1/e}}{2 \log_2 q} \cdot 8^{-r} r^{-1} (\log_2 s)^{1-1/r}. \qquad (2)\end{aligned}$$

Let $G \in \mathbb{G}(M_s^r, p)$. By Lemma 1, with some suitable probability π computed hereafter, $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(H_t^r)$ for every δ. From Lemma 3, $\mathrm{IRS}_\delta(H_t^r) = \Omega((\delta+2)^{1-r} r^{-2} 2^{-r}(t-7)^{r-1})$ for every δ. By the choice of t, and by inequations (2), $\mathrm{IRS}_\delta(H_t^r) = \Omega(16^{-r}(\delta+2)^{1-r} r^{-3} (\log s)^{1-1/r})$.

Let us compute the probability π. From Lemma 1, $\pi \geqslant 1 - \exp(-mq^{dx})$, where d is the maximum degree of the base graph (here $B = M_s^r$), x is the number of nodes of the m-subgraph (here $x = |V(H_t^r)| = t^r$) and recall that q is a non-null constant $\leqslant \frac{1}{2}$. We have $x = \alpha r \log_2 s$, $d = 2r$, and $m = \lfloor s/(t+2) \rfloor^r$. For s large enough, $m \geqslant (s/(2t))^r$.

We have $q^{dx} = q^{2r^2 \alpha \log_2(s/2)} = (s/2)^{-1/4}$. We have also that $(s/(2t))^r = (s/2)^r/t^r = (s/2)^r/(\alpha r \log_2 s)$. We observe that $\alpha r = -1/(8r \log_2 q) \leqslant c$, where c is some constant since $r \geqslant 1$ and q is a constant. So, for s large enough it follows that

$$m \geqslant (s/(2t))^r \geqslant (s/2)^r/(c \log_2 s) \geqslant (s/2)^{r-1/4}$$

and so $mq^{dx} \geqslant (s/2)^{r-1/2}$. Therefore,

$$\pi \geqslant 1 - \exp\left(-(s/2)^{r-1/2}\right)$$

as claimed, that completes the proof. □

Corollary 1. *For all integers $r \in [1, \log_2 n]$ and $\delta \geqslant 0$, with probability at least $1 - \exp(-\sqrt{n})$, a random r-dimensional mesh G of n nodes with constant failure probability verifies that* $\mathrm{IRS}_\delta(G) = \Omega(16^{-r}(\delta+2)^{1-r} r^{-4} (\log n)^{1-1/r})$.

Proof. The number of nodes in M_s^r is $n = s^r$. We observe that for $r = 1$ or $s = \mathrm{O}(1)$, the result holds since in both cases we have that $16^{-r}(\delta+2)^{1-r} r^{-4} (\log n)^{1-1/r} = \mathrm{O}(1)$ and $\mathrm{IRS}_\delta(G)$ is always at least equal to 1 (i.e., with probability 1). We note that for $r \geqslant 2$ and for s large enough, we have $(s/2)^{r-1/2} \geqslant s^{r/2} = \sqrt{n}$. □

[2] W.l.o.g. α is chosen such that t is integer. Indeed, if $\alpha = -1/(8r^2 \log_2 q)$ does not suffice, we can enlarge α by a factor at most two. This does not affect the final result.

3. Upper bound for random two-dimensional meshes

It is also of great importance to state a nontrivial upper bound on IRS_δ for random r-dimensional meshes, and in particular the upper bound on IRS_0, that is the compactness of a random mesh. The trivial upper bound is $\mathrm{O}(n)$. In this section we consider random meshes $\mathbb{G}(M_s^2, p)$ with $n = s^2$ nodes and for $p = 0.5$, that is the critical value in the percolation theory.

We will use the following property due to [14].

Lemma 4 (*Gavoille and Guévremont [14]*). *The compactness of any connected graph is the maximum of the compactness over all its biconnected components.*

Note that the compactness of an n-node graph is no more that $n/2$. Actually, from [17], the compactness is no more than $n/4+\mathrm{o}(n)$, and there are graphs with compactness $n/4-\mathrm{o}(n)$. So, in general, the compactness of G is at most $\mathrm{O}(L)$, where L is the size of the largest biconnected components of G.

Let $f(s)$ denote the expected value of the number of nodes in the largest biconnected component in a random $s \times s$ mesh (with $p = 0.5$). In order to estimate $f(s)$ on large values of s we have programmed a standard linear-time algorithm for finding biconnected components based on the modification of a depth first search traversal, cf. [36]. Some outputs of our algorithm are depicted in Figs. 5–9 (see also Fig. 4).

Our experiments (up to 50 tests for random meshes as large as 2048×2048 meshes) show that $f(s) \approx \beta s^\alpha$ where $\beta = 0.32$ and $\alpha = 1.654$. (More precisely, we remark that

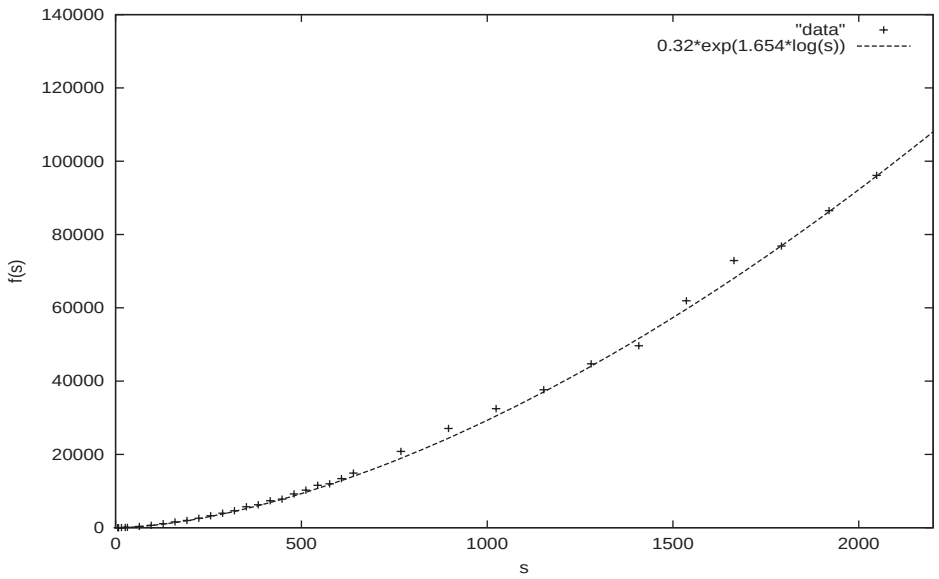

Fig. 4. The size of the largest biconnected components of a random $s \times s$ mesh. The tests come from averaging 2000 experiments (for the small values of s) and up to 50 experiments (for the larger values of s).

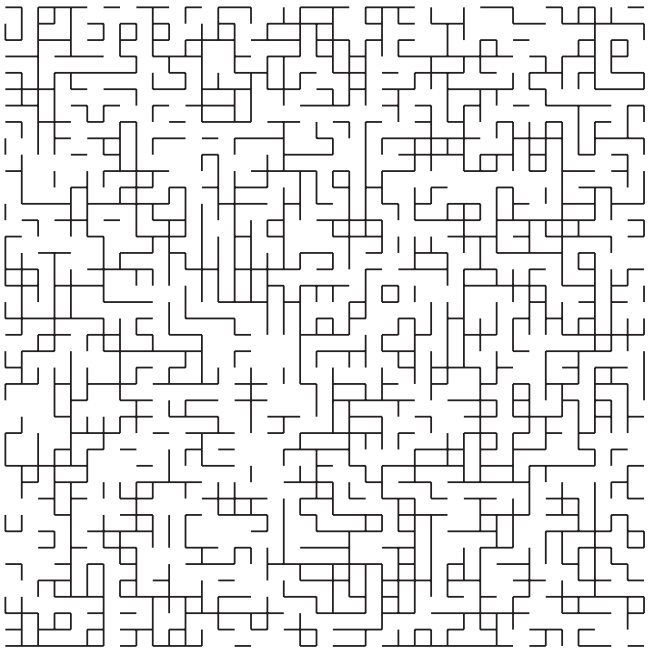

Fig. 5. A 40×40 random mesh with $p = 0.5$. Edges are horizontal and vertical segments, isolated nodes have been removed.

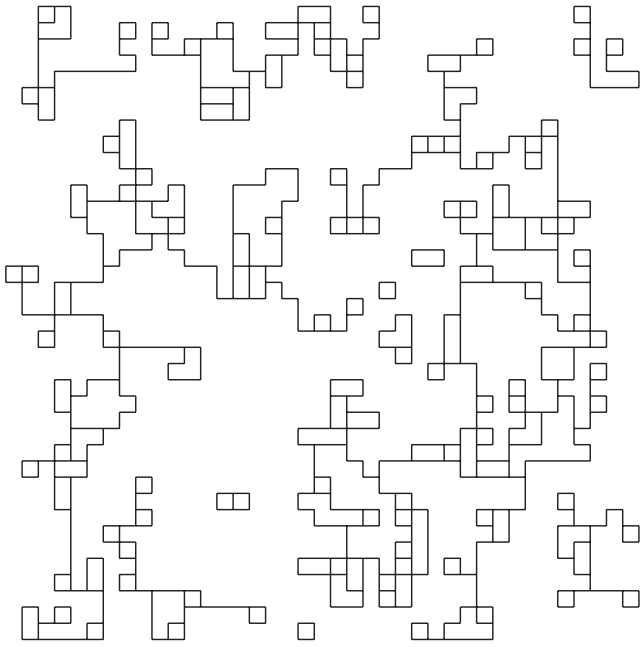

Fig. 6. The mesh of Fig. 5 without trees.

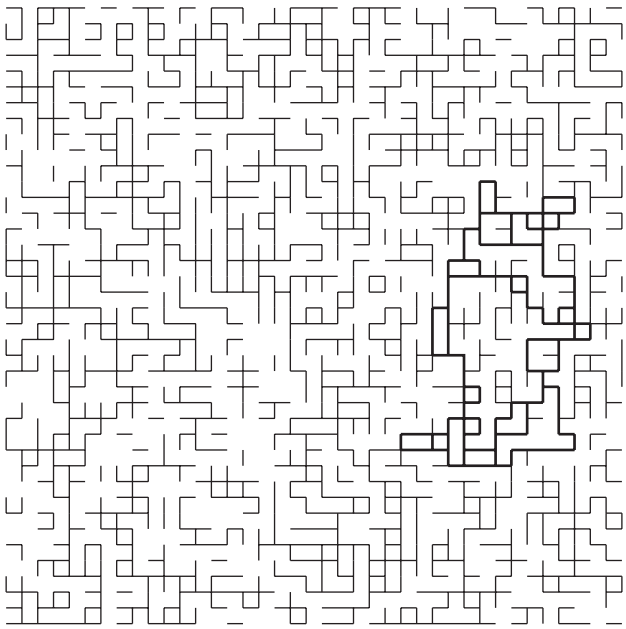

Fig. 7. The largest biconnected component (with 116 nodes) in Fig. 5.

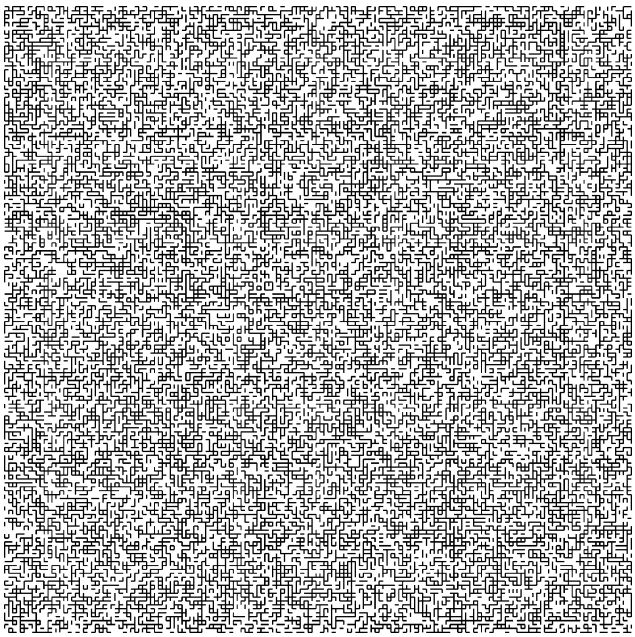

Fig. 8. A 160×160 random mesh with $p = 0.5$.

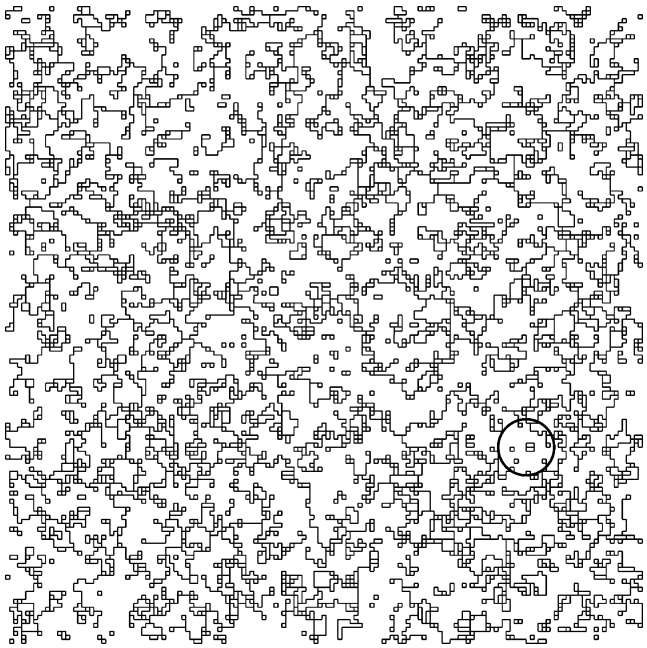

Fig. 9. The mesh of Fig. 8 without trees. The largest biconnected component has 672 nodes. Observe that this random mesh contains the H_3^2 graph (circled) as isolated subgraph.

$f(2x)/f(x)$ is a constant, so $\log f(x) \approx ax + b$.) The value α has been computed by $\alpha = \log_2 A$, where A is the average value of $f(2^{i+1})/f(2^i)$ for $i = 1, \ldots, 10$. The value on β follows. Expressed as the total number of nodes of the mesh ($n = s^2$), it turns out that $f(s) \approx \Theta(n^{0.827})$. Our experimental results are summarized in Fig. 4.

4. Conclusion and further works

We have proved a nonconstant asymptotic lower bound on the number of intervals for δ-stretched interval routing schemes on random multi-dimensional meshes. By the experimental simulations we have also stated a nontrivial upper bound on this measure for the two-dimensional case.

We leave several open questions and further directions for the study of compact routing in reliability networks.

1. Extension to arbitrary routing strategies. It would be interesting to prove similar lower bounds in a general encoding model, so applicable to any encoding of the routing scheme rather than the interval routing model.
2. Extension to random B-graphs with different edge probabilities. For instance a model $\mathbb{G}(B, p_r, p_c)$ where B is a mesh and where p_r applies to rows and p_c to the columns would be interesting. In particular, finding the shortest-path routing tables complexity of

$G \in \mathbb{G}(B, 0.5, 1)$ is of some interest. This could also be a tool for the study of augmented random graphs for Small World (cf. Kleingberg's model).

3. We know that w.h.p. random K_n-graphs (the $\mathbb{G}(n, p)$ model) have constant compactness [3] [18], and that random meshes have compactness $\Omega(16^{-r} r^{-4} (\log n)^{1-1/r})$ (this paper). Complete graphs and meshes having compactness 1. However, can we have a base graph B of high compactness, say $\Omega(n^{\varepsilon})$, such that random B-graphs have low compactness (w.h.p.)?
4. What is the compactness of a general n-node random B-graph for $p = 0.5$ (still w.h.p.)? Is $\Omega(n)$ possible? Same question if B is a bounded degree graph or a planar graph or if B is the hypercube (our lower bound on r-dimensional meshes just gives a constant in this case).
5. Is H_t^2 the worst-case sub-mesh for the compactness? If the answer is "yes", then it would clearly improve our $O(n^{0.827})$ upper bound for random meshes. They would have compactness at most $O(\sqrt{n})$.
6. But maybe, the most interesting question remains: is there an n-node planar graph with compactness larger than $\sqrt{n}$?

Acknowledgements

The authors are grateful to the anonymous referees for their helpful and constructive comments.

References

[1] I. Benjamini, O. Schramm, Percolation beyond Z^d: many questions and a few answers, Electron. Commun. Probab. 1 (1996) 71–82.
[2] B. Bollobás, Random Graphs, second ed., Cambridge Studies in Advanced Mathmatics, Vol. 73, 2001.
[3] L.J. Cowen, Compact routing with minimum stretch, J. Algorithms 38 (2001) 170–183.
[4] Y. Dourisboure, C. Gavoille, Improved compact routing scheme for chordal graphs, in: 16th Internat. Symp. on Distributed Computing (DISC), Lecture Notes in Computer Science, Vol. 2508, Springer, Berlin, October 2002, pp. 252–264.
[5] T. Eilam, C. Gavoille, D. Peleg, Compact routing schemes with low stretch factor, J. Algorithms 46 (2003) 97–114.
[6] T. Eilam, S. Moran, S. Zaks, The complexity of the characterization of networks supporting shortest-path interval routing, in: Fourth Internat. Coll. on Structural Information & Communication Complexity (SIROCCO), Carleton Scientific, July 1997, pp. 99–111.
[7] M. Elkin, D. Peleg, $(1+\varepsilon, \beta)$-spanner constructions for general graphs, in: 33rd Ann. ACM Symp. on Theory of Computing (STOC), Hersonissos, Crete, Greece, July 2001, pp. 173–182.
[8] M. Flammini, J. van Leeuwen, A.M. Spaccamela, The complexity of interval routing on random graphs, Comput. J. 41 (1) (1998) 16–25.
[9] P. Fraigniaud, C. Gavoille, Universal routing schemes, J. Distributed Comput. 10 (1997) 65–78.
[11] G.N. Frederickson, R. Janardan, Efficient message routing in planar networks, SIAM J. Comput. 18 (4) (1989) 843–857.
[12] C. Gavoille, A survey on interval routing, Theoret. Comput. Sci. 245 (2) (2000) 217–253.
[13] C. Gavoille, Routing in distributed networks: overview and open problems, ACM SIGACT News—Distributed Comput. Column 32 (1) (2001) 36–52.

[3] That is the minimum number of intervals for shortest-path interval routing schemes.

[14] C. Gavoille, E. Guévremont, Worst case bounds for shortest path interval routing, J. Algorithms 27 (1998) 1–25.
[15] C. Gavoille, N. Hanusse, Compact routing tables for graphs of bounded genus, in: 26th Internat. Coll. on Automata, Languages and Programming (ICALP), Lecture Notes in Computer Science, Vol. 1644, Springer, Berlin, July 1999, pp. 351–360.
[16] C. Gavoille, M. Katz, N.A. Katz, C. Paul, D. Peleg, Approximate distance labeling schemes, in: Nineth Ann. European Symp. on Algorithms (ESA), Lecture Notes in Computer Science, Vol. 2161, Springer, Berlin, August 2001, pp. 476–488.
[17] C. Gavoille, D. Peleg, The compactness of interval routing, SIAM J. Discrete Math. 12 (4) (1999) 459–473.
[18] C. Gavoille, D. Peleg, The compactness of interval routing for almost all graphs, SIAM J. Comput. 31 (3) (2001) 706–721.
[19] C. Gavoille, S. Pérennès, Lower bounds for interval routing on 3-regular networks, in: Third Internat. Coll. on Structural Informations & Communication Complexity (SIROCCO), Carleton University Press, June 1996, pp. 88–103.
[20] C. Gavoille, D. Peleg, S. Pérennes, R. Raz, Distance labeling in graphs, in: 12th Symp. on Discrete Algorithms (SODA), ACM-SIAM, January 2001, pp. 210–219.
[21] G.R. Grimmett, Percolation, Springer, Berlin, May 1999.
[22] S. Janson, T. Łuczak, A. Ruciński, Random Graphs, Wiley, New York, 2000.
[23] J. Kleinberg, The small-world phenomenon: an algorithmic perspective, in: 32nd Ann. ACM Symp. on Theory of Computing (STOC), May 2000, pp. 163–170.
[24] J. Kleinberg, Navigation in a small world, Nature 406 (2000) 845.
[25] R. Kráľovič, P. Ružička, D. Štefankovič, The complexity of shortest path and dilation bounded interval routing, Theoret. Comput. Sci. 234 (1–2) (2000) 85–107.
[26] H.-I. Lu, Improved compact routing tables for planar networks via orderly spanning trees, in: Eighth Ann. Internat. Computing & Combinatorics Conf. (COCOON), Lecture Notes in Computer Science, Vol. 2387, Springer, Berlin, August 2002, pp. 57–66.
[27] M. Nehéz, The compactness lower bound of shortest-path interval routing on $n \times n$ tori with random faulty links, Tech. Report 582, KAM-DIMATIA Series, Charles University, Praha, 2002.
[28] S.E. Nikoletseas, K.V. Palem, P.G. Spirakis, M. Yung, Short vertex disjoint paths and multiconnectivity in random graphs: reliable network computing, in: 21st Internat. Coll. on Automata, Languages and Programming (ICALP), Lecture Notes in Computer Science, Vol. 820, Springer, Berlin, 1994, pp. 508–515 (preliminary version of [29]).
[29] S.E. Nikoletseas, K.V. Palem, P.G. Spirakis, M. Yung, Connectivity properties in random regular graphs with edge faults, Internat. J. Foundations of Computer Science (IJFCS) 11 (2) (2000) 247–262.
[30] S.E. Nikoletseas, P.G. Spirakis, Expander properties in random regular graphs with edge faults, in: 12th Ann. Symp. on Theoretical Aspects of Computer Science (STACS), Lecture Notes in Computer Science, Vol. 900, Springer, Berlin, 1995, pp. 421–432.
[31] D. Peleg, Proximity-preserving labeling schemes, J. Graph Theory 33 (2000) 167–176.
[32] D. Peleg, E. Upfal, A trade-off between space and efficiency for routing tables, J. ACM 36 (3) (1989) 510–530.
[33] N. Robertson, P.D. Seymour, Graph minors. III. planar tree-width, J. Combin. Theory 36 (1984) 49–64.
[34] P. Ružička, On efficiency of path systems induced by routing and communication schemes, Comput. Inform. 20 (2001) 181–205.
[35] N. Santoro, R. Khatib, Labelling and implicit routing in networks, Comput. J. 28 (1) (1985) 5–8.
[36] R.E. Tarjan, Depth-first search and linear graph algorithms, SIAM J. Comput. 1 (2) (1972) 146–160.
[37] M. Thorup, Compact oracles for reachability and approximate distances in planar digraphs, in: 42th Ann. IEEE Symp. on Foundations of Computer Science (FOCS), IEEE Computer Society Press, Silverspring, MD, October 2001.
[38] M. Thorup, U. Zwick, Compact routing schemes, in: 13th Ann. ACM Symp. on Parallel Algorithms and Architectures (SPAA), Hersonissos, Crete, Greece, July 2001, ACM PRESS, New York, pp. 1–10.
[39] S.S.H. Tse, F.C.M. Lau, On the space requirement of interval routing, IEEE Trans. Comput. 48 (7) (1999) 752–757.
[40] J. van Leeuwen, R.B. Tan, Interval routing, Comput. J. 30 (4) (1987) 298–307.