Laboratoire Bordelais de Recherche en Informatique
UMR 5800 - Université Bordeaux I, 351, cours de la Libération,
33405 Talence CEDEX, France

Research Report RR-1294-03

Interval Routing in Reliability Networks

by Cyril Gavoille and Martin Nehéz

March, 2003

Interval Routing in Reliability Networks

Cyril Gavoille[*]   Martin Nehéz [†]

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 destroying each edge of B 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 with $p=0.5$, that is the percolation model at the critical phase. We show a lower bound of $\Omega(\sqrt{\log n})$ on the number of intervals required per edge for every additive stretch $\delta \leqslant c\sqrt{\log n}$ for some constant $c>0$. On the other side, our experiments results show that the size of the largest biconnected components is at most $O(n^{0.827})$, and thus that there exists a trivial shortest-path routing scheme using $O(n^{0.827})$ intervals per edge. The results are extended to random meshes of dimension higher than two.

Keywords: compact routing tables, reliability networks, random graphs, interval routing, percolation theory

1 Introduction

Research in the area of routing algorithms on computer networks became of great interest by many researchers in several recent years. This topic is interesting both from the theoretical and also from the application point of view. One aspect of research in this field is to compact 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

[*]Laboratoire Bordelais de Recherche en Informatique, Université Bordeaux I. *gavoille@labri.fr*

[†]Dept. of Computer Science and Engineering, Slovak University of Technology, Ilkovičova 3, 812 19 Bratislava. *nehez@elf.stuba.sk* Supported by the VEGA grant VG 1/0162/03.

and their distance in the network. Among them [Cow01, EGP02, FG97, PU89, TZ01] and [Gav01] for a survey.

The above strategies apply to all the networks, however is it natural to wonder whether other more efficient techniques can be applied on realistic network. 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 [Kle00b, Kle00a]).

1.1 Reliability Networks

In that 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 destroying each edge of B 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 [NPSY94, NPSY00] and is described in more details in [JŁR00, pp. 2]. Reliability network is a natural generalization of the binomial random graph model of Erdös-Rényi, denoted hereafter $\mathbb{G}(n, p)$, for which $B = K_n$ the complete graph on n nodes. The reliability network based on infinite square mesh belongs to percolation theory [Gri99]. As mentioned in [JŁR00], this model can be generalized further by allowing different probabilities of failure at different edges. This model is also related to the other problems of computer science such as grid-computing, fault-tolerance 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 [EP01] and distance computation [GKK$^+$01], and yields also compact routing tables [DG02].

1.2 Routing Tables and Interval Routing

Finally we focus on *interval routing schemes*, a particular way of implementing standard routing tables [SK85, vLT87]. Recall that a routing table on a graph G consists of the set of distinct addresses ranging in $\{1, \ldots, |V(G)|\}$ and a set of *local routing tables* associated

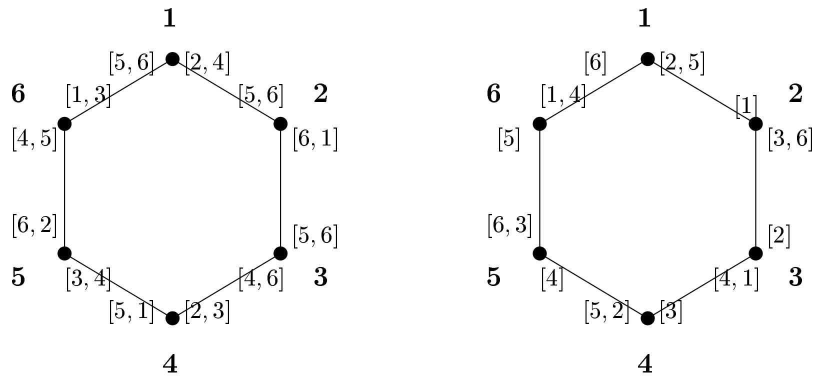

Figure 1: A shortest-path 1-IRS (left side), and a 2-stretched 1-IRS (right side).

with each node. 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 the message through the output port number q which is computed by u by looking at the ith entry of its local table. So the route is computed in a distributed fashion by all 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 table as follows: u stores a d entry tables, d the degree of u. The qth entry is the list of destination address v for which the route from u to v uses output port number q. If for all the nodes, the destination addresses using the same output port can be grouped into 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 [Gav00] and [Ruž01].

The main difficulty in the design of interval routing scheme 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 [EMZ97]. 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 $O(n \log d)$ bits are required for a standard routing table implementation. 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 supports shortest-path $O(1)$-IRS (more results about interval routing is accessible in the survey [Gav00]).

Flammini, van Leeuwen and Marchetti-Spaccamela [FvLMS98] proved a non-constant lower bound on k for shortest-path k-IRS on random graph 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 $p = n^{-1+1/\Theta\sqrt{\log n}}$. On the other hand, Gavoille and Peleg [GP01] proved that almost all graphs (that is graph 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 one having 1 interval, leaving open the question of whether almost graphs support shortest-path 1-IRS. Finally, for shortest-path k-IRS on random n-node tori, a preliminary result appears in [Neh02] where it is proved a lower bound of $\Omega(\sqrt{\log n})$.

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 mesh. We show that asymptotically almost surely, $k = \Omega(\delta^{1-r} r^{-3} (\log n)^{1-1/r})$, for every additive stretch $\delta \leqslant c(\log n)^{1/r}$ for some constant $c > 0$.
2. For upper bounds, we have studied random square meshes 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 components [Gri99]. Unfortunately, as many interesting problems in Percolation Theory (and as suggested by [BS96]), we are reduced to make experiments. Based on the expected size of the largest biconnected components of a 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 is a typical planar graph (cf. the Graph Minor Theory of Robertson and Seymour [RS84]). 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 $\Omega(\sqrt{n})$ [EGP02] and $O(n)$ bits per node [GH99, Lu02]. For shortest-path interval routing, the range is similar: $\Omega(\sqrt{n})$ intervals is the best known lower bound [GP96, TL99], and $O(n)$ is the trivial upper bound. Similar gaps exist also for *distance labeling* in planar graphs whose goal here is to compute distances between two nodes based only on a node label [Pel00]: 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 [GPPR01]. Finding shortest-paths and distances structure of 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 [FJ89, Tho01]. 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 by a constant δ. This forces any $O(\delta)$-stretched routing scheme to respect the shortest paths in the subdivided graph of that remains of linear size.

The paper is organized as follows: in 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 a graph. Because the graphs of $\mathbb{G}(B,p)$ are not necessarily connected, we extend the notion of routing scheme on non-connected graphs as follows: a routing scheme on a non-connected graph is simply the union of the routing scheme on 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 G_i's are the connected components of G.

For two integers $r \geqslant 1$ and $s \geqslant 2$, the *r-dimensional mesh*, denoted M_s^r, is the graph whose the nodes are all the r-tuples over the set $\{0,\dots,s-1\}$. Two nodes $u = (u_0,\dots,u_{r-1})$ and $v = (v_0,\dots,v_{r-1})$, 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 our lower bound result we need several preliminary results.

Let X be an induced subgraph of graph G. X is said to be a *subgraph of δ-stretched paths* of G if for all $u,v \in V(X)$, every δ-stretched path between u and v in G is wholly contained in X. For instance the subgraph induced by 3 consecutive nodes of the 6-cycle is a subgraph of 2-stretched paths (cf. Fig. 1). The following property is a generalization of the result of [FG98], originally stated for $\delta = 0$.

Lemma 1 *If X is a subgraph of δ-stretched paths of G, then* $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(X)$.

Proof. Let X be a subgraph of δ-stretched paths of G. Let R be a (k,δ)-interval routing scheme for G such that $k = \mathrm{IRS}_\delta(G)$. If X is a subgraph of δ-stretched paths of G, then all δ-stretched paths between any pair of nodes of X are wholly contained in X. Hence R is a (k,δ)-interval routing scheme for X with node labels taken from the set $\{1,\dots,|V(G)|\}$. It suffices to replace each node label by its rank in the range $\{1,\dots,|V(X)|\}$ to obtain an interval routing scheme R' on X. R' is a (k,δ)-IRS as the scheme R (see details in Theorem 4 of [FG98], p. 167). It yields that $\mathrm{IRS}_\delta(X) \leqslant k$. $\square$

A subgraph X of G is *isolated* if there is no edge $\{u,v\} \in E(G)$ such that $u \in V(X)$ and $v \notin V(X)$. X is *near-isolated* if there is only one edge $\{u,v\} \in E(G)$ such that $u \in V(X)$ and $v \notin V(X)$ (this edge must be a bridge in G). It is clear that any simple path connecting two nodes taken in an isolated or near-isolated subgraph X is wholly contained in X, in particular any simple path that is not a shortest path. So, any isolated or near-isolated subgraph is a subgraph of δ-stretched paths in G, and this for every δ. By Lemma 1, it follows that for every connected graph G that contains a near-isolated subgraph X, $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(X)$, for every δ.

A graph X is a *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 on Fig. 1.

Lemma 2 *Let B a connected graph with n nodes and maximum degree d, and let X be a m-subgraph of B with x nodes. Then, $G \in \mathbb{G}(B,p)$ contains X as subgraph of δ-stretched paths with probability at least $1-\exp(-m\,q^{dx})$, where $q = \min\{p, 1-p\}$. So, by Lemma 1, $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(X)$, for every δ.*

Proof. Let $G \in \mathbb{G}(B,p)$, and let A be the event "G contains X as subgraph of δ-stretched paths". Our goal is to lower bound $\Pr(A)$. From the previous discussion, if X is near-isolated (or isolated) in G, then G contains X as subgraph of δ-stretched paths. So,

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

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 $Z = \sum_{i=1}^m Z_i$. The result now comes easily from the Chernoff's bound [HR90].

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

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 near-isolated in G" is an event independent from "X_j is near-isolated in G". The variables Z_i are mutually independent. It follows that,

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

To make X_i isolated or near-isolated in B it suffices to keep or destroy independently some edges of B whose at least one extremity is in X_i. We 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 $(1-b/a)^c \leqslant e^{-(b/a)c}$, for all $0 < b \leqslant b+c < a$ (cf. [Bol01, Eq. (1.6) pp. 5]). This completes the proof. □

In order to obtain lower bounds on the number of intervals we use a result of Královič, Ružička and Štefankovič [KRŠ00]. For every arc e incident to a node $v \in V$, let $S(v, e)$ denote the subset of nodes $w \in V$ such that at least one shortest-path from v to w traverses e.

Lemma 3 ([KRŠ00]) *Let G be a graph with maximum degree d. Let Q and W be two disjoint node subsets of G such that for all distinct nodes $w_i, w_j \in W$, there is $v \in Q$ such that for each incident arc e of v it holds $w_i \notin S(v,e)$ or $w_j \notin S(v,e)$. Then, every $(k,0)$-IRS on G must satisfy*

$$k \geqslant \frac{|W|}{d|Q|}.$$

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 every $t, r \geqslant 0$, we define the graph H_t^r recursively by $H_t^r = H_t^{r-1} \star P_t$, with $H_t^0 = P_1$ the graph with one node. 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 subgraph of the r-dimensional mesh M_s^r for $t \leqslant s$. The number of nodes of H_t^r is t^r. It holds the following recurrent equation for the number of edges: $|E(H_t^{i+1})| = 2|E(H_t^i)| + (t-1)t^i$ with $|E(H_t^0)| = 0$. Fig. 2 depicts the H_t^2 graph.

Lemma 4 $\mathrm{IRS}_\delta(H_t^r) = \Omega\left(\delta^{1-r} r^{-2} t^{r-1}\right)$*, for all $\delta \leqslant t/3 - 2$ and $r \geqslant 2$.*

Proof. We will extend the property of the sets Q and W from Lemma 3 to (k,δ)-interval routing scheme for the graph H_t^r. We will proceed by induction.

1. Let $r = 2$. The graph H_t^2 consists of t^2 nodes which are arranged into t *rows* each of t nodes. Nodes of only the two rows (the topmost and the bottommost) are connected by the path. On the other hand, nodes of each *column* is connected by the path. According to Lemma 3 we will construct the sets Q^2 and W^2. Let us fix $\delta \geqslant 0$. Let the set Q^2 contain all the nodes which lie on the bottommost row and on the leftmost column. Let the set W^2 consist of row nodes except the nodes from Q^2 as follows: the first row which belongs to W^2 has the horizontal distance from the bottommost row $\delta + 2$. The next row of W has the same spacing from the previous one, etc. The last row which belongs into W^2 has the distance at least $\delta/2 + 1$ from the topmost row of H_t^2. (See Fig. 2.) We check that the property for the sets Q^2 and W^2 of Lemma 3 and its extension to (k,δ)-IRS holds. The argument is as follows.

(a) If nodes $w_i, w_j \in W^2$, $w_i \neq w_j$, are from the same column of H_t^2 then the corresponding node $q \in Q^2$ lies on the leftmost column of H_t^2 such that the difference of the lengths between the nodes w_i, q and w_j, q, respectively, is at most 1. (It depends on the fact whether t is even or odd.) For example, for w_1 and w_t the corresponding node from Q^2 is $q_{\lceil 3\delta/2\rceil+3}$. Node q is incident with two edges, say e and e'. A message sent from q and addressed to w_i can be delivered to w_i by two possible ways: via the arc which corresponds to the edge e or via the arc corresponding to e'. If the length of the first routing path is l then the length of the second one must be at least $l+\delta+1$ or vice versa. It means that the only one possibility mentioned satisfies the property of δ-stretched routing. On the other hand, the opposite situation holds if the message is sent from q towards w_j. Such the way, for each arc e outgoing from q, both w_i, w_j cannot be reached from q over its outgoing arc e using δ-stretched paths.

(b) If nodes $w_i, w_j \in W^2$ are from the different columns then the corresponding node $q \in Q^2$ lies on the bottommost row of H_t^2. According to the requirement of δ-stretched routing, this situation is analogous to the previous one.

Thus, for the sets Q^2 and W^2 as above, for all the pairs $w_i, w_j \in W^2$, $w_i \neq w_j$, there is $q \in Q^2$ such that for each arc e outgoing from q, both w_i, w_j cannot be reached from q over its outgoing arc e using δ-stretched paths for a given δ. For $\delta = 0$ it is the optimal routing and hence, by Lemma 3: $\mathrm{IRS}_0(H_t^2) \geqslant |W^2|/(3\cdot|Q^2|)$, since $|Q^2| = 2t-1$, $|W^2| \geqslant (t-1)\,(t/(\delta+2)-2)$ and $\delta \leqslant t/3-2$. From the construction of the sets Q^2 and W^2 follows that each (k,δ)-interval routing scheme for H_t^2 must have the same path system with respect to Q^2 and W^2 as a given shortest-path interval routing scheme for Q^2 and W^2. It follows that $\mathrm{IRS}_\delta(H_t^2) \geqslant \mathrm{IRS}_0(H_t^k) = \Omega(t/\delta)$ for $\delta \leqslant t/3-2$.

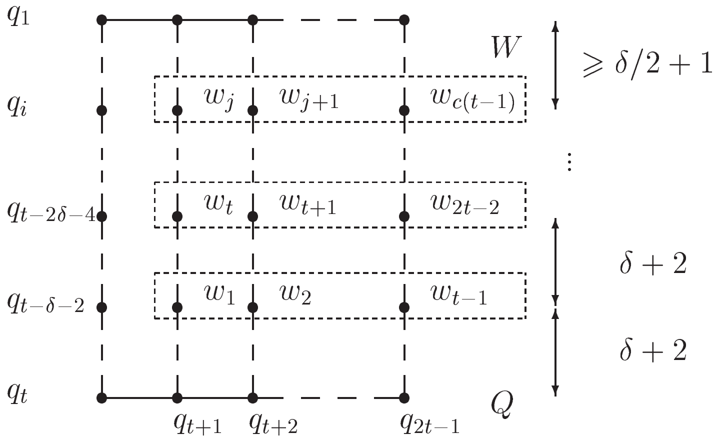

Figure 2: The sets W^2 and Q^2 in the graph H_t^2.

2. Let $r > 2$. The set Q^r contains all nodes from Q^{r-1} and also all nodes from the path in rth dimension which is crossing graph H_t^{r-1} in the node q_t. The set W^r is constructed by the same manner as in the case $r=2$. The set W^{r-1} consists of at least $(t-1)\,(t/(\delta+2)-2)$ nodes layered in the $(r-1)$-dimensional hyperplane $\mathcal{H}^{r-1}$. The set W^r will be consists of at least $(t-\delta-2)/(\delta+2)-1$ hyperplanes $\mathcal{H}^{r-1}$ (without edges) such that each $\mathcal{H}^{r-1}$ will be orthogonal to the paths oriented along rth dimension. The

distance between two hyperplanes $\mathcal{H}^{r-1}$ in rth dimension is $\delta+2$ and the distance of the last hyperplane from the top is at least $\delta/2+1$, similarly as in the case $r=2$. We will show that the sets Q^r and W^r satisfy the condition of the Lemma 3 extended to the δ-stretched routing. As in the previous situation, we will distinguish two cases.

(a) If the nodes $w_i, w_j \in W^r$, $w_i \neq w_j$, are from the same path then the corresponding node $q \in Q^r$ lies on the parallel path (oriented along the rth dimension) which belongs to Q^r. By the same argument as for $r=2$ it holds that for all the pairs $w_i, w_j \in W^r$, $w_i \neq w_j$, there is $q \in Q^r$ such that for each arc e outgoing from q, both w_i, w_j cannot be reached from q over its outgoing arc e using δ-stretched paths.

(b) If the nodes $w_i, w_j \in W^r$, $w_i \neq w_j$, are from the different paths (the paths are along the rth dimension) then we take the projection of w_i and w_j into the hyperplane $\mathcal{H}^{r-1}$ which contains the set Q^{r-1} and that is orthogonal to any path in rth dimension. By this projection we obtain the images w_i^{r-1}, w_j^{r-1} of the nodes w_i, w_j such that $w_i^{r-1}, w_j^{r-1} \in H_t^{r-1}$. From the recursive construction of W^r follows that the nodes w_i^{r-1}, w_j^{r-1} lies in the set W^{r-1} so that there exists a node $q \in Q^{r-1}$ with the property that for each arc e outgoing from q, both w_i^{r-1}, w_j^{r-1} cannot be reached from q over its outgoing arc e using δ-stretched paths. Thus, the node q is also the same for the nodes $w_i, w_j \in W^r$ because the distances between w_i and w_i^{r-1}, or between w_j and w_j^{r-1}, respectively, do not affect to the position of q.

By Lemma 3, if the routing in H_t^r is shortest-path then, according to the construction of the sets Q^r and W^r it holds $\mathrm{IRS}_0(H_t^r) = \Omega\left(\frac{t^{r-1}}{\delta^{r-1}r^2}\right)$, since $|Q^r| = rt-1$, $|W^r| \geqslant (t-1)\left(t/(\delta+2)-2\right)^{r-1}$ and $\delta \leqslant t/3-2$. On the other hand, from the construction of the sets Q^r and W^r follows that each (k,δ)-interval routing scheme for H_t^r must have the same path system with respect to Q^r and W^r as a given shortest-path interval routing scheme for Q^r and W^r. It follows that $\mathrm{IRS}_\delta(H_t^r) \geqslant \mathrm{IRS}_0(H_t^r) = \Omega\left(\frac{t^{r-1}}{\delta^{r-1}r^2}\right)$ for $\delta \leqslant t/3-2$. $\square$

Remark. It is not difficult to construct a (k,δ)-IRS for H_t^2 with $k = O(t/\delta)$ intervals. Therefore, for $r=2$, the bound of Lemma 4 is tight, i.e., $\mathrm{IRS}_\delta(H_t^2) = \Theta(t/\delta)$.

The main result of this section is the following.

Theorem 1 *Let $0 < p < 1$ be a constant and let r, s be integers such that $s^r \to +\infty$, $r \geqslant 2$, and s large enough. With probability at least $1 - \exp(-\sqrt{s}\,)$ a random mesh $G \in \mathbb{G}(M_s^r, p)$ verifies that $\mathrm{IRS}_\delta(G) = \Omega(\delta^{1-r} r^{-3} (\log s)^{1-1/r})$, for every $\delta \leqslant c(\log s)^{1/r}$, for some constant $c > 0$.*

Proof. Let us show that H_t^r is a m-subgraph of M_s^r for $m = \lfloor s/(t+2) \rfloor^r$ and for every $t \leqslant s-2$. 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 M_{t+2}^r. Each sub-mesh contains H_t^r as a subgraph one can centered so that any two copies of H_t^r into M_{t+2}^r in 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^r_t and pairwise at distance two. For $t \leqslant s-2$, $m \geqslant 1$, and H^r_t is a m-subgraph of M^r_s as claimed.

Let us fix $t = (\alpha r \log s)^{1/r}$ for some $\alpha = \Theta(1/r^2)$ given later. Clearly, for s large enough and $s^r \to +\infty$, $t \leqslant s-2$. Let $G \in \mathbb{G}(M^r_s, p)$. By Lemma 2, with some suitable probability π computed hereafter, $\mathrm{IRS}_\delta(G) \geqslant \mathrm{IRS}_\delta(H^r_t)$ for every δ. From Lemma 4, $\mathrm{IRS}_\delta(H^r_t) = \Omega(\delta^{1-r} r^{-2} t^{r-1})$ for every $\delta \leqslant t/3-2$. By the choose of t,

$$t^{r-1} \;=\; (\alpha r \log s)^{1-1/r} \;=\; \Theta(r^{-1}(\log s)^{1-1/r})$$

as $r^{O(1/r)} = \Theta(1)$ for every $r \geqslant 2$. Hence $\mathrm{IRS}_\delta(H^r_t) \geqslant \Omega(\delta^{1-r} r^{-3} (\log s)^{1-1/r})$ for every $\delta \leqslant t/3-2$. Let $\beta = \frac{1}{4}(\alpha r)^{1/r}$. We have $\beta(\log s)^{1/r} \leqslant t/3-2$ as $s^r \to +\infty$. For $\alpha = \Theta(1/r^2)$, $\beta = \Theta(1)$, so β is lower bounded by a constant $c > 0$, and $c(\log s)^{1/r} \leqslant t/3-2$. It follows, that for every $\delta \leqslant c(\log s)^{1/r}$, $\mathrm{IRS}_\delta(H^r_t) \geqslant \Omega(\delta^{1-r} r^{-3} (\log s)^{1-1/r})$.

Let us compute the probability π. From Lemma 2, $\pi = 1 - \exp(-mq^{dx})$ where $q = \min\{p, 1-p\}$ is a constant, d is a maximum degree of the base graph (here M^r_t), and where x the number of nodes of the m-subgraph (here $x = |V(H^r_t)| = t^r$). We have $x = t^r = \alpha r \log s$, $d = 2r$, and $m = \lfloor s/(t+2) \rfloor = \Theta(s/(\log s)^{1/r})$. Plugging all previous values, we obtain:

$$\pi \;\geqslant\; 1 - \exp\left(-\Theta\left(\frac{s}{(\log s)^{1/r}} \cdot q^{2r^2 \alpha \log s}\right)\right) .$$

Let us fix $\alpha = -1/(8r^2 \log q)$. As q is a constant $\leqslant 1/2$, we have $\alpha = \Theta(1/r^2)$ as required. So, $q^{2r^2\alpha \log s} = s^{-1/4}$. We have also $s/(\log s)^{1/r} \geqslant s^{3/4}$. Therefore,

$$\pi \;\geqslant\; 1 - \exp\left(-\sqrt{s}\right)$$

as claimed, that completes the proof. □

Observing that $n = s^r$ and that $r^{1/r} \in [1, \sqrt{2}]$ for every $r \geqslant 2$, we have also:

Corollary 1 *For every $\delta \leqslant c\sqrt{\log n}$ for some constant $c > 0$, with probability at least $1 - \exp(-n^{1/(2r)})$, a random r-dimensional mesh G of n nodes with constant failure probability verifies that* $\mathrm{IRS}_\delta(G) = \Omega(\delta^{1-r} r^{-3} (\log n)^{1-1/r})$.

And for $r = 2$,

Corollary 2 *For every $\delta \leqslant c\sqrt{\log n}$ for some constant $c > 0$, with probability at least $1 - \exp(-n^{1/4})$, a random 2-dimensional mesh G of n nodes verifies that* $\mathrm{IRS}_\delta(G) = \Omega(\sqrt{\log n})$.

3 Upper Bound for Random 2-dimensional Meshes

This is also of great importance to state the nontrivial upper bound on IRS_δ for random r-dimensional meshes, and in particular upper bound on IRS_0, the compactness of a random mesh. The trivial upper bound is $O(n)$. 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 [GG98].

Lemma 5 ([GG98]) *Let G be a connected graph. The compactness of any graph G is the maximum of the compactness overall biconnected components of G.*

Note that the compactness of an n-node graph is no more that $n/2$. Actually, from [GP99], the compactness is no more than $n/4 + o(n)$ if $n/\log n \to +\infty$, and there are graphs with compactness $n/4 - o(n)$. So, in general, the compactness of G is at most $O(L)$, where L is the size of the largest biconnected components of G.

Let $f(s)$ denotes the excepted 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. [Tar72]. Some outputs of our algorithm are depicted on figures 4, 5, 6, 7, and 8 (in Appendix).

Our experiments (up to 50 tests for random meshes as large as 2048×2048 meshes) show that $f(s) \approx \alpha s^\beta$ where $\alpha = 0.32$ and $\beta = 1.654$. The value β has been computed by $\beta = \log_2 A$, where A is the average value of $f(2^{i+1})/f(2^i)$ for $i = 1, \ldots, 10$. The value on α follows. Expressed as the total number of nodes of the mesh ($n = s^2$), it turns out that $f(s) \approx O(n^{0.827})$. Our experimental results are summarized on Fig. 3.

4 Conclusion and Further Works

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

1. Extension to arbitrary routing strategy. 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 probability. 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 find the shortest-path routing table complexity of $G \in \mathbb{G}(B, 0.5, 1)$ is of particular interest. This could also be a tool for the study of augmented random graphs for Small World (cf. Kleingberg's model).

3. Can we have a base graph B of high compactness[1] but with low compactness of random B-graph (w.h.p.)? At least two examples state the reverse. We know that ran-

[1] That is the number of intervals for shortest-path interval routing scheme.

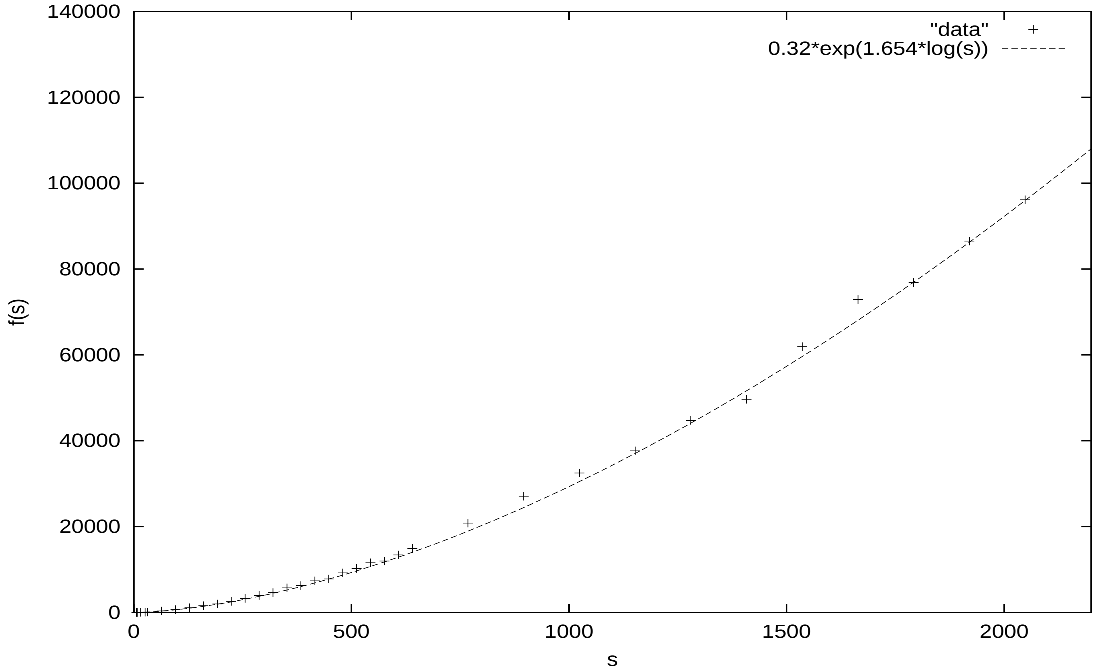

Figure 3: The size of the largest biconnected components of a random $s \times s$ mesh. The tests come from averaging 2000 experiments (for the small values of s) and up to 50 experiments (for the larger values of s).

dom K_n-graphs (the $\mathbb{G}(n,p)$ model) have constant compactness [GP01], and that random meshes have compactness $\Omega(\sqrt{\log n})$ (this paper), complete graphs and meshes having compactness 1.

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 give a constant is 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 $\sqrt{n}$?

References

[Bol01] Béla Bollobás. *Random Graphs (2nd edition)*. Cambridge Studies in Advanced Mathmatics 73, 2001.

[BS96] Itai Benjamini and Oded Schramm. Percolation beyond Z^d: many questions and a few answers. *Electronic Communications in Probability*, 1:71–82, 1996.

[Cow01] Lenore J. Cowen. Compact routing with minimum stretch. *Journal of Algorithms*, 38:170–183, 2001.

[DG02] Yon Dourisboure and Cyril Gavoille. Improved compact routing scheme for chordal graphs. In 16^{th} *International Symposium on Distributed Computing (DISC)*, volume 2508 of Lecture Notes in Computer Science, pages 252–264. Springer, October 2002.

[EGP02] Tamar Eilam, Cyril Gavoille, and David Peleg. Compact routing schemes with low stretch factor. *Journal of Algorithms*, 2002. To appear.

[EMZ97] Tamar Eilam, Shlomo Moran, and Shmuel Zaks. The complexity of the characterization of networks supporting shortest-path interval routing. In Danny Krizanc and Peter Widmayer, editors, 4^{th} *International Colloquium on Structural Information & Communication Complexity (SIROCCO)*, pages 99–111. Carleton Scientific, July 1997.

[EP01] Michael Elkin and David Peleg. $(1+\epsilon,\beta)$-spanner constructions for general graphs. In 33^{rd} *Annual ACM Symposium on Theory of Computing (STOC)*, pages 173–182, Hersonissos, Crete, Greece, July 2001.

[FG97] Pierre Fraigniaud and Cyril Gavoille. Universal routing schemes. *Journal of Distributed Computing*, 10:65–78, 1997.

[FG98] Pierre Fraigniaud and Cyril Gavoille. Interval routing schemes. *Algorithmica*, 21:155–182, 1998.

[FJ89] Greg N. Frederickson and Ravi Janardan. Efficient message routing in planar networks. *SIAM Journal on Computing*, 18(4):843–857, August 1989.

[FvLMS98] Michele Flammini, Jan van Leeuwen, and Alberto Marchetti-Spaccamela. The complexity of interval routing on random graphs. *The Computer Journal*, 41(1):16–25, 1998.

[Gav00] Cyril Gavoille. A survey on interval routing. *Theoretical Computer Science*, 245(2):217–253, 2000.

[Gav01] Cyril Gavoille. Routing in distributed networks: Overview and open problems. *ACM SIGACT News - Distributed Computing Column*, 32(1):36–52, March 2001.

[GG98] Cyril Gavoille and Eric Guévremont. Worst case bounds for shortest path interval routing. *Journal of Algorithms*, 27:1–25, 1998.

[GH99] Cyril Gavoille and Nicolas Hanusse. Compact routing tables for graphs of bounded genus. In Jiří Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, 26^{th} *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 1644 of Lecture Notes in Computer Science, pages 351–360. Springer, July 1999.

[GKK$^+$01] Cyril Gavoille, Michal Katz, Nir A. Katz, Christophe Paul, and David Peleg. Approximate distance labeling schemes. In F. Meyer auf der Heide, editor, 9^{th} *Annual European Symposium on Algorithms (ESA)*, volume 2161 of Lecture Notes in Computer Science, pages 476–488. Springer, August 2001.

[GP96] Cyril Gavoille and Stéphane Pérennès. Lower bounds for interval routing on 3-regular networks. In Nicola Santoro and Paul Spirakis, editors, 3^{rd} *International Colloquium on Structural Information & Communication Complexity (SIROCCO)*, pages 88–103. Carleton University Press, June 1996.

[GP99] Cyril Gavoille and David Peleg. The compactness of interval routing. *SIAM Journal on Discrete Mathematics*, 12(4):459–473, October 1999.

[GP01] Cyril Gavoille and David Peleg. The compactness of interval routing for almost all graphs. *SIAM Journal on Computing*, 31(3):706–721, 2001.

[GPPR01] Cyril Gavoille, David Peleg, Stéphane Pérennes, and Ran Raz. Distance labeling in graphs. In 12^{th} *Symposium on Discrete Algorithms (SODA)*, pages 210–219. ACM-SIAM, January 2001.

[Gri99] Geoffrey R. Grimmett. *Percolation.* Springer, May 1999.

[HR90] Torben Hagerup and Christine Rüb. A guided tour of chernoff bounds. *Information Processing Letters*, 33:305–308, 1990.

[JŁR00] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs.* Wiley-Interscience Publication, 2000.

[Kle00a] Jon Kleinberg. Navigation in a small world. *Nature*, 406:845, August 2000.

[Kle00b] Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In 32^{nd} *Annual ACM Symposium on Theory of Computing (STOC)*, pages 163–170, May 2000.

[KRŠ00] Rastislav Kráľovič, Peter Ružička, and Daniel Štefankovič. The complexity of shortest path and dilation bounded interval routing. *Theoretical Computer Science*, 234(1-2):85–107, 2000.

[Lu02] Hsueh-I Lu. Improved compact routing tables for planar networks via orderly spanning trees. In 8^{th} *Annual International Computing & Combinatorics Conference (COCOON)*, volume 2387 of Lectures Notes in Computer Science, pages 57–66. Springer, August 2002.

[Neh02] Martin Nehéz. The compactness lower bound of shortest-path interval routing on $n \times n$ tori with random faulty links. Technical Report 582, KAM-DIMATIA Series, Charles University, Praha, 2002.

[NPSY94] S.E. Nikoletseas, K.V. Palem, Paul G. Spirakis, and Moti Yung. Short vertex disjoint paths and multiconnectivity in random graphs: Reliable network computing. In 21^{st} *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 820 of Lecture Notes in Computer Science, pages 508–515. Springer-Verlag, 1994. (preliminary version of [NPSY00]).

[NPSY00] S.E. Nikoletseas, K.V. Palem, Paul G. Spirakis, and Moti Yung. Connectivity properties in random regular graphs with edge faults. *International Journal of Foundations of Computer Science (IJFCS)*, 11(2):247–262, 2000.

[Pel00] David Peleg. Proximity-preserving labeling schemes. *Journal of Graph Theory*, 33:167–176, 2000.

[PU89] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36(3):510–530, July 1989.

[RS84] Neil Robertson and Paul D. Seymour. Graph minors. III. planar tree-width. *Journal of Combinatorial Theory*, 36:49–64, 1984.

[Ruž01] Peter Ružička. On efficiency of path systems induced by routing and communication schemes. *Computing and Informatics*, 20:181–205, 2001.

[SK85] Nicola Santoro and Ramez Khatib. Labelling and implicit routing in networks. *The Computer Journal*, 28(1):5–8, February 1985.

[Tar72] Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, June 1972.

[Tho01] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. In 42^{th} *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society Press, October 2001.

[TL99] Savio S. H. Tse and Francis C. M. Lau. On the space requirement of interval routing. *IEEE Transactions on Computers*, 48(7):752–757, July 1999.

[TZ01] Mikkel Thorup and Uri Zwick. Compact routing schemes. In 13^{th} *Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 1–10, Hersonissos, Crete, Greece, July 2001. ACM PRESS.

[vLT87] Jan van Leeuwen and Richard B. Tan. Interval routing. *The Computer Journal*, 30(4):298–307, 1987.

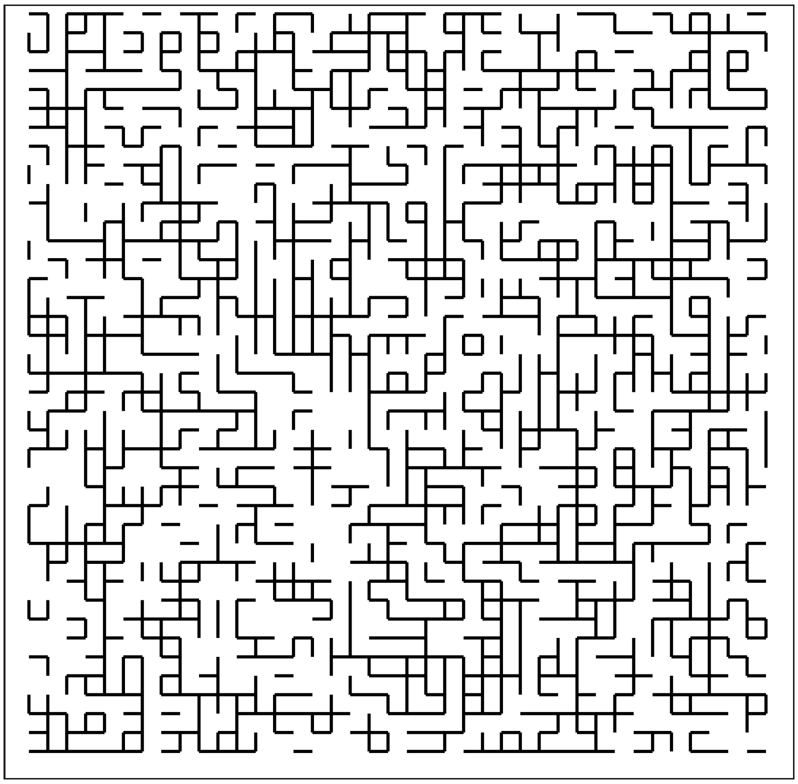

Figure 4: A 40×40 random mesh with $p = 0.5$. Edges are horizontal and vertical segments, isolated nodes have been removed.

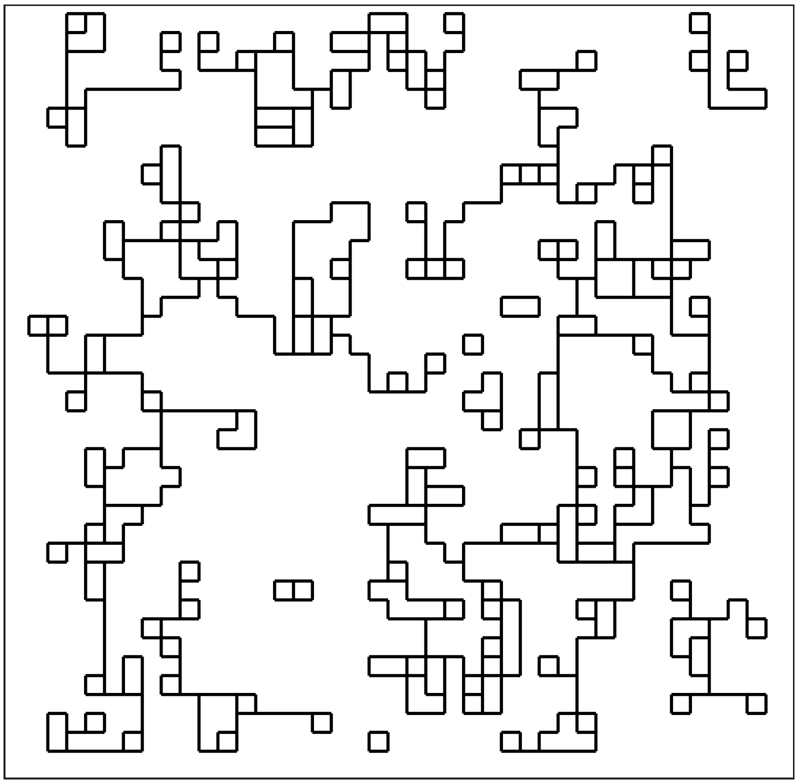

Figure 5: The mesh of Fig. 4 without trees.

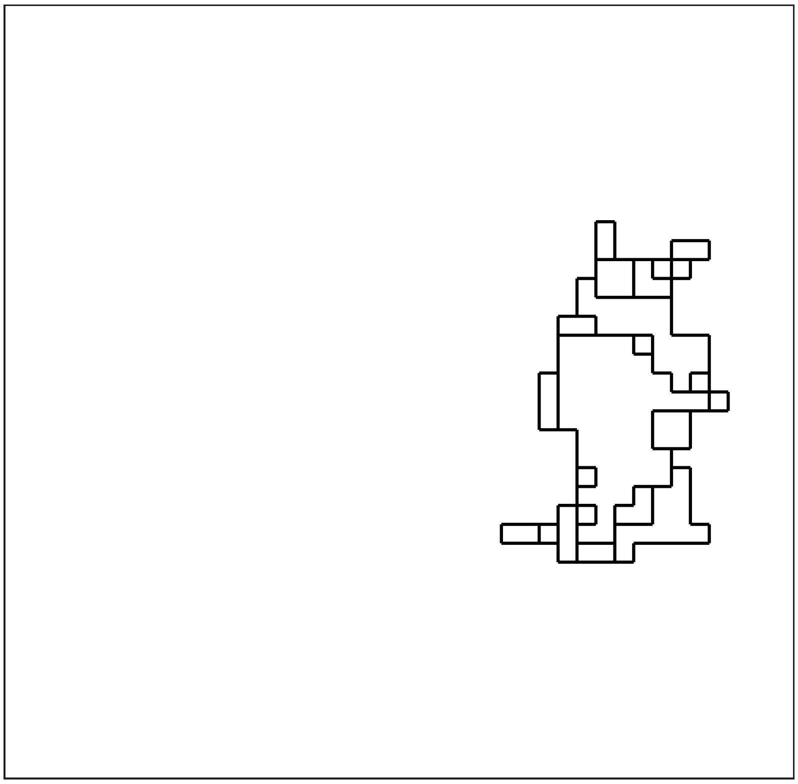

Figure 6: The largest biconnected component (with 116 nodes) in Fig. 4.

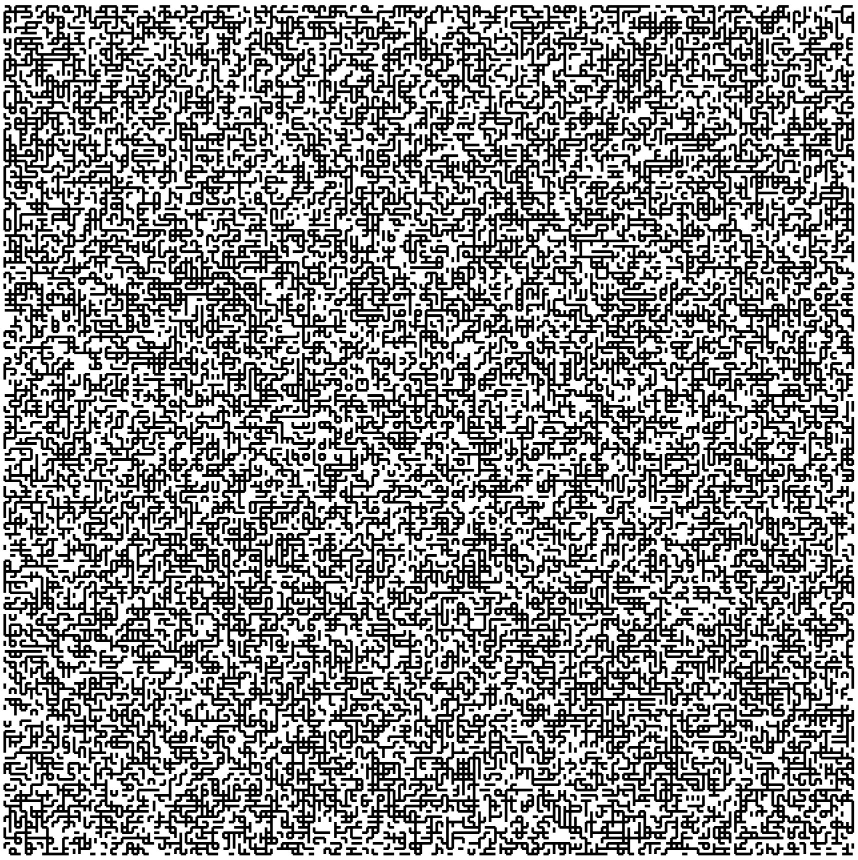

Figure 7: A 160×160 random mesh with $p = 0.5$.

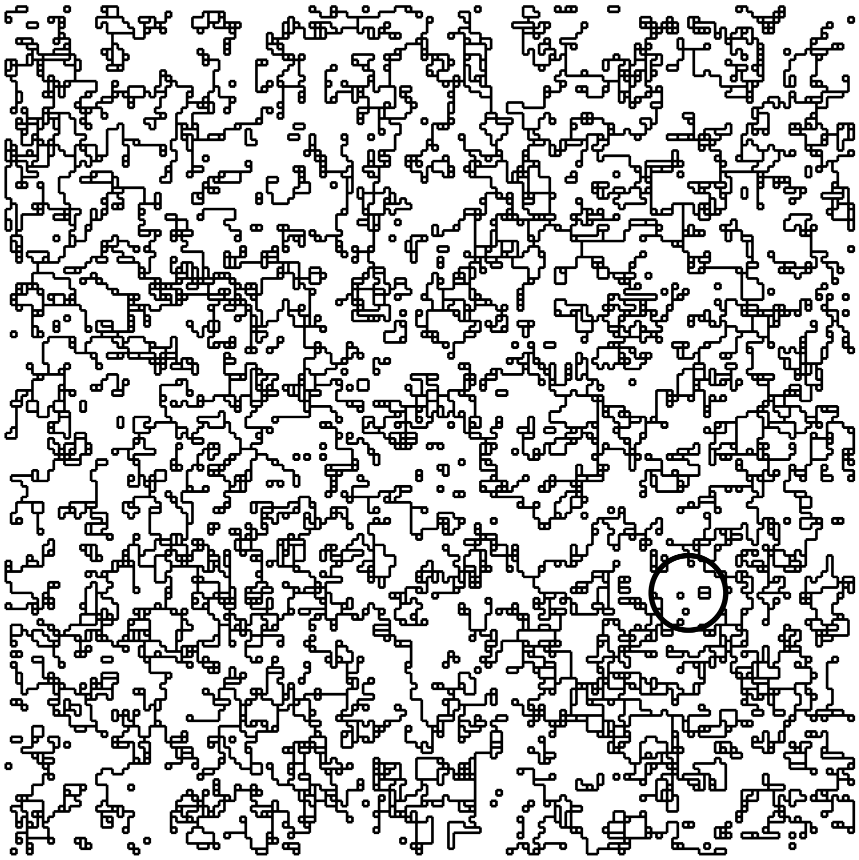

Figure 8: The mesh of Fig. 7 without trees. The largest biconnected component has 672 nodes. Observe that this random mesh contains the H_3^2 graph (circled) as isolated subgraph.