

Interval Routing in Reliability Networks*

CYRIL GAVOILLE

Université Bordeaux I, France

MARTIN NEHÉZ

Slovak University of Technology, Bratislava

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 $\Omega(\sqrt{\log n}/(\delta + 1))$ on the number of intervals required per edge for every additive stretch $\delta \geq 0$. 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 $\Omega((\delta + 1)^{1-r} r^{-4} (\log n)^{1-1/r})$, for every additive stretch $\delta \geq 0$ and for every integral dimension $r \in [1, \log_2 n]$.

Keywords

compact routing tables, reliability networks, random graphs, percolation theory

1 Introduction

Research in the area of routing algorithms on computer networks became of great interest by many researchers in 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.

*Supported by the VEGA grant VG 1/0162/03.

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, 30, 35] and [12] 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 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, 22]).

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 $\{1, \dots, |V(B)|\}$, 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 [27, 28] and is described in more detail in [21, pp. 2]. The reliability network is a natural generalization of the binomial random graph model of Erdős-Rényi, denoted hereafter $\mathbb{G}(n, p)$, for which $B = K_n$ is the complete graph on n nodes. The reliability network based on infinite square mesh belongs to percolation theory [20]. As mentioned in [21], this model can be generalized further by allowing different probabilities of failure at different edges. This model 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 [15], and yields also compact routing tables [4].

1.2 Routing Tables and Interval Routing

We focus on *interval routing scheme*, a particular way of implementing standard routing tables [33, 37]. Recall that a routing table on a graph G consists of the set of distinct addresses ranging in $\{1, \dots, |V(G)|\}$ and a set of *local routing tables* associated 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 i th 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 tables as follows: u stores a d entry tables, d being the degree of u . The q th entry is the list of destination addresses 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 [11] and [32].

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 minimum while keeping the routes near-shortest paths. Whenever shortest paths are required, the problem to decide 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 $O(n \log d)$ 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 support shortest-path $O(1)$ -IRS (more results about interval routing are accessible in the survey [11]).

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 Erdős-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 [17] proved that almost all graphs (that is 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




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

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 all graphs support shortest-path 1-IRS. Finally, for shortest-path k -IRS on random n -node tori, a preliminary result appears in [26] 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 meshes with constant failure probability. We show that asymptotically almost surely, $k = \Omega((\delta + 1)^{1-r} r^{-4} (\log n)^{1-1/r})$, for every additive stretch $\delta \geq 0$ 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 components [20]. Unfortunately, as many interesting problems in Percolation Theory (and as suggested by [1]), 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 graphs (cf. the Graph Minor Theory of Robertson and Sey-

mour [31]). 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 [14, 25]. For shortest-path interval routing, the range is similar: $\Omega(\sqrt{n})$ intervals is the best known lower bound [19, 36], 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 their node label [29]: 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 [18]. Finding structure of shortest paths and distances of planar graphs is probably difficult and certainly would require more combinatorics.

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 $\text{IRS}_\delta(G)$ the smallest integer k such that G supports a (k, δ) -IRS. The number $\text{IRS}_0(G)$ 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 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, $\text{IRS}_\delta(G) = \max_i \text{IRS}_\delta(G_i)$ where the G_i 's are the connected components of G .

For two integers $r \geq 1$ and $s \geq 1$, the r -dimensional mesh, denoted M_s^r , is the graph whose nodes are all the r -tuples over the set $\{1, \dots, s\}$. Two nodes $u = (u_1, \dots, u_r)$ and $v = (v_1, \dots, v_r)$ 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.

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 contained in X . For instance the subgraph induced by 3 consecutive nodes of the 6-cycle is a subgraph of 1-stretched paths (cf. Fig. 1). The following property is a generalization of the result of [10], originally stated for $\delta = 0$.

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

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 $k = \text{IRS}_\delta(G)$. If X is a subgraph of δ -stretched

paths of G , then all δ -stretched paths between any pair of nodes of X are 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 [10], p. 167). It implies that $\text{IRS}_\delta(X) \leq k$. \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 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 , $\text{IRS}_\delta(G) \geq \text{IRS}_\delta(X)$, 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 (take two opposite edges).

Lemma 2 *Let B be 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 $1 - \exp(-mq^{dx})$, where $q = \min\{p, 1 - p\}$. So, by Lemma 1, $\text{IRS}_\delta(G) \geq \text{IRS}_\delta(X)$, 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,

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

Let X_1, \dots, X_m be m subgraphs of B pairwise at distance two, each one isomorphic to X . For each $i \in \{1, \dots, m\}$, 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 $Z = \sum_{i=1}^m Z_i$.

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

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,

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

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 $\Pr(Z_i = 1) \geq q^{dx}$, where $q = \min\{p, 1 - p\}$. Thus,

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

using the fact that $(1 - b/a)^c \leq e^{-(b/a)c}$, for all $0 < b \leq b + c < a$ (cf. [2, Eq. (1.6) pp. 5]). This completes the proof. \square

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 R on G has the property Π (or is a Π -routing) if all the routes induced by R belongs 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, Ružička and Štefankovič [24] originally proved for the shortest-path property.

Lemma 3 *Let G be a graph, and Π be a routing property. Let $P(u, v)$ 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 $w \notin P(u, v)$ or $w' \notin P(u, v)$. Then, every k -IRS with property Π on G must satisfy*

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

Proof. Let R be any k -IRS with property Π on G , a graph with node set $\{v_1, \dots, v_n\}$ and arc set (each edge appears twice, one for each orientation) $\{e_1, \dots, e_m\}$. 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 j th column of $M_{i,j}$ (the last bit and the first bit being considered as consecutive). Let $c(e_j)$ be its 01-sequence number. From the choice of R , $c(e_j) \leq k$ 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-sequence number of the column e_j of M' . Removing some bits of a binary vector does not increase its 01-sequence number. Hence, $c'(e_j) \leq c(e_j)$, and so $k \geq c'(e_j)$.

Let us show that $\sum_i c'(e_j) \geq |W|$, where the sum is done over all the columns of M' . Indeed, consider two consecutive rows of M' (again the last and the first

column 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 , $w \notin P(u, v)$ or $w' \notin P(u, v)$. R has the property Π , thus u has a neighbor v' such that $w' \in P(u, v')$. As $w \notin P(u, v')$, it follows that the binary vector of the column associated with the arc (u, v') contains a 01-sequence starting at position 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., $\sum_i c'(e_j) \geq |W|$ as claimed.

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

$$k \geq \frac{\sum_i c'(e_j)}{\sum_{u \in U} \deg(u)} \geq \frac{|W|}{\sum_{u \in U} \deg(u)} .$$

□

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 $t, r \geq 1$, 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,

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

See Fig. 2 for an example. Note that H_t^r is a subgraph of the r -dimensional mesh M_s^r for $t \leq s$.

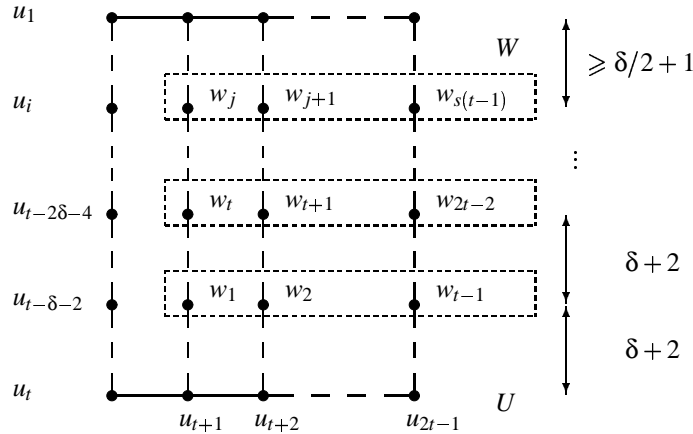


Figure 2: The sets W and U in the graph H_t^2 .

Lemma 4 $\text{IRS}_\delta(H_t^r) = \Omega((\delta + 1)^{1-r} r^{-2} t^{r-1})$, for all $\delta \geq 0$ and $r \geq 1$.

Proof Sketch. The result holds trivially for $r = 1$ or for $\delta > t/3 - 2$. So assume $r \geq 2$ and $\delta \leq t/3 - 2$.

Due to space limitation, we only sketch the proof for $r = 2$. The sets $U = \{u_i\}$ and $W = \{w_i\}$ are chosen as depicted on Fig. 2. We check that $|W| \geq (t - 1)(t/(\delta + 2) - 2)^{r-1}$ and $\sum_{u \in U^r} \deg(u) \leq (r + 1)(rt - 1)$, so that $|W| / \sum_{u \in U} \deg(u) = \Omega((\delta + 1)^{1-r} r^{-2} t^{r-1})$. Then the result comes from Lemma 3 applied to the routing property $\Pi = \text{“}\delta\text{-stretched path”}$. \square

Remark: It is not difficult to construct a (k, δ) -IRS for H_t^2 with $k = O(t/(\delta + 1))$ intervals. Therefore, for $r = 2$, the bound of Lemma 4 is tight, i.e., $\text{IRS}_\delta(H_t^2) = \Theta(t/(\delta + 1))$.

The main result of this section is the following.

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

Proof. The result holds for $r = 1$, as $(\delta + 1)^{1-r} r^{-3} (\log s)^{1-1/r} = 1$ and as every $G \in \mathbb{G}(M_s^1, p)$ is a forest thus satisfying $\text{IRS}_0(G) = 1$ (cf. [33]). We note also that for $s = O(1)$, $(\delta + 1)^{1-r} r^{-3} (\log s)^{1-1/r} = O(1)$, thus the result holds as well. So let us assume $r \geq 2$ and that s is large enough (i.e., s is greater than some fixed constant s_0).

Let us show that, for every $t \leq s - 2$, H_t^r is a 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 that 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_t^r and pairwise at distance two. For $t \leq s - 2$, $m \geq 1$ and H_t^r is a m -subgraph of M_s^r as claimed.

Let us fix $t = (\alpha r \log s)^{1/r}$ for some $\alpha = \Theta(1/r^2)$ given later. Note that $r^{1/r} \in [1, e^{1/e}]$ for every $r \geq 1$. Thus, $(\alpha r)^{1/r} = \Theta(1)$ and $t = \Theta((\log s)^{1/r})$. So, for s large enough and $r \geq 2$, $t \leq s - 2$. Let $G \in \mathbb{G}(M_s^r, p)$. By Lemma 2, with some suitable probability π computed hereafter, $\text{IRS}_\delta(G) \geq \text{IRS}_\delta(H_t^r)$ for every δ . From Lemma 4, $\text{IRS}_\delta(H_t^r) = \Omega((\delta + 1)^{1-r} r^{-2} t^{r-1})$ for every δ . By the choice of t ,

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

Hence $\text{IRS}_\delta(H_t^r) = \Omega((\delta + 1)^{1-r} r^{-3} (\log s)^{1-1/r})$.

Let us compute the probability π . From Lemma 2, $\pi \geq 1 - \exp(-mq^{dx})$ where $q = \min\{p, 1 - p\}$ is a constant, d is the maximum degree of the base graph (here $B = M_t^r$), and where x is the number of nodes of the m -subgraph (here $x =$

$|V(H_t^r)| = t^r$). We have $x = \alpha r \log s$, $d = 2r$, and $m = \lfloor s/(t+2) \rfloor^r \geq (s/(3t))^r = \Theta(r(s/3)^r / \log s)$. Plugging all previous values, we obtain:

$$\pi \geq 1 - \exp\left(-\Theta\left(\frac{r(s/3)^r}{\log s} \cdot q^{2r^2 \alpha \log s}\right)\right).$$

Let us fix $\alpha = 1/(8r^2 \log q)$. As q is a non-null constant $\leq 1/2$, we have $\alpha = \Theta(1/r^2)$ as required. So, $q^{2r^2 \alpha \log s} = s^{-1/4}$. We have also, for s large enough and $r \geq 2$, that $\Theta(r(s/3)^r / \log s) \geq (s/3)^{r-1/4}$. Therefore,

$$\pi \geq 1 - \exp\left(- (s/3)^{r-1/2}\right)$$

as claimed, that completes the proof. \square

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

Proof. The number of nodes in M_s^r is $n = s^r$. We observe that for $r = 1$ or $s = O(1)$, the result holds since in both cases we have that $(\delta+1)^{1-r} r^{-4} (\log n)^{1-1/r} = O(1)$ and $\text{IRS}_\delta(G) \geq 1$ always (with probability 1). We note that for $r \geq 2$ and for s large enough, we have $(s/3)^{r-1/2} \geq s^{r/2} = \sqrt{n}$. \square

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 [13].

Lemma 5 ([13]) *The compactness of any connected graph is the maximum of the compactness overall its biconnected components.*

Note that the compactness of an n -node graph is no more than $n/2$. Actually, from [16], the compactness is no more than $n/4 + o(n)$, 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 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. [34]. Some outputs of our algorithm are depicted on figures 4, and 5.

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 $f(2x)/f(x)$ 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 $i = 1, \dots, 10$. 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 on Fig. 3.




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).

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 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.




Figure 4: A 40×40 random mesh with $p = 0.5$. Edges are horizontal and vertical segments, isolated nodes have been removed. The largest biconnected component, in bold, has 116 nodes.

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 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. We know that w.h.p. random K_n -graphs (the $\mathbb{G}(n, p)$ model) have constant compactness* [17], and that random meshes have compactness $\Omega(\sqrt{\log n})$ (this paper). Complete graphs and meshes having compactness 1. However, can we have a base graph B of high compactness, say $\Omega(n^\epsilon)$, 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

*That is the minimum number of intervals for shortest-path interval routing scheme.




Figure 5: A 160×160 random mesh with $p = 0.5$, depicted without trees (i.e., nodes of degree ≤ 1 have been successively removed). The largest biconnected component has 672 nodes. Observe that this random mesh contains the H_3^2 graph (circled) as isolated subgraph.

graph? or a planar graph? or if B is the hypercube (our lower bound on r -dimensional meshes just give 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} ?

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

References

- [1] BENJAMINI, I., AND SCHRAMM, O. Percolation beyond Z^d : many questions and a few answers. *Electronic Communications in Probability* 1 (1996), 71–82.
- [2] BOLLOBÁS, B. *Random Graphs (2nd edition)*. Cambridge Studies in Advanced Mathematics 73, 2001.
- [3] COWEN, L. J. Compact routing with minimum stretch. *Journal of Algorithms* 38 (2001), 170–183.
- [4] DOURISBOURE, Y., AND GAVOILLE, C. Improved compact routing scheme for chordal graphs. In *DISC (2002)*, vol. 2508 of LNCS, pp. 252–264.
- [5] EILAM, T., GAVOILLE, C., AND PELEG, D. Compact routing schemes with low stretch factor. *Journal of Algorithms* 46 (2003), 97–114.
- [6] EILAM, T., MORAN, S., AND ZAKS, S. The complexity of the characterization of networks supporting shortest-path interval routing. In *SIROCCO (1997)*, Carleton Scientific, pp. 99–111.
- [7] ELKIN, M., AND PELEG, D. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. In *STOC (2001)*, pp. 173–182.
- [8] FLAMMINI, M., VAN LEEUWEN, J., AND MARCHETTI-SPACCAMELA, A. The complexity of interval routing on random graphs. *The Computer Journal* 41, 1 (1998), 16–25.
- [9] FRAIGNIAUD, P., AND GAVOILLE, C. Universal routing schemes. *Journal of Distributed Computing* 10 (1997), 65–78.
- [10] FRAIGNIAUD, P., AND GAVOILLE, C. Interval routing schemes. *Algorithmica* 21 (1998), 155–182.
- [11] GAVOILLE, C. A survey on interval routing. *Theoretical Computer Science* 245, 2 (2000), 217–253.
- [12] GAVOILLE, C. Routing in distributed networks: Overview and open problems. *ACM SIGACT News - Distrib. Comp. Column* 32, 1 (2001), 36–52.
- [13] GAVOILLE, C., AND GUÉVREMONT, E. Worst case bounds for shortest path interval routing. *Journal of Algorithms* 27 (1998), 1–25.

-
- [14] GAVOILLE, C., AND HANUSSE, N. Compact routing tables for graphs of bounded genus. In *ICALP (1999)*, vol. 1644 of LNCS, pp. 351–360.
- [15] GAVOILLE, C., KATZ, M., KATZ, N. A., PAUL, C., AND PELEG, D. Approximate distance labeling schemes. In *ESA (2001)*, vol. 2161 of LNCS, Springer, pp. 476–488.
- [16] GAVOILLE, C., AND PELEG, D. The compactness of interval routing. *SIAM Journal on Discrete Mathematics* 12, 4 (1999), 459–473.
- [17] GAVOILLE, C., AND PELEG, D. The compactness of interval routing for almost all graphs. *SIAM Journal on Computing* 31, 3 (2001), 706–721.
- [18] GAVOILLE, C., PELEG, D., PÉRENNES, S., AND RAZ, R. Distance labeling in graphs. In *SODA (2001)*, ACM-SIAM, pp. 210–219.
- [19] GAVOILLE, C., AND PÉRENNÈS, S. Lower bounds for interval routing on 3-regular networks. In *SIROCCO (1996)*, Carleton Univ. Press, pp. 88–103.
- [20] GRIMMETT, G. R. *Percolation*. Springer, 1999.
- [21] JANSON, S., LUCZAK, T., AND RUCIŃSKI, A. *Random Graphs*. Wiley-Interscience Publication, 2000.
- [22] KLEINBERG, J. Navigation in a small world. *Nature* 406 (2000), 845.
- [23] KLEINBERG, J. The small-world phenomenon: An algorithmic perspective. In *STOC (2000)*, pp. 163–170.
- [24] KRÁĽOVIČ, R., RUŽIČKA, P., AND ŠTEFANKOVIČ, D. The complexity of shortest path and dilation bounded interval routing. *Theoretical Computer Science* 234, 1-2 (2000), 85–107.
- [25] LU, H.-I. Improved compact routing tables for planar networks via orderly spanning trees. In *COCOON (2002)*, vol. 2387 of LNCS, pp. 57–66.
- [26] NEHÉZ, M. The compactness lower bound of shortest-path interval routing on $n \times n$ tori with random faulty links. Tech. Rep. 582, KAM-DIMATIA Series, Charles University, Praha, 2002.
- [27] NIKOLETSEAS, S., PALEM, K., SPIRAKIS, P. G., AND YUNG, M. Short vertex disjoint paths and multiconnectivity in random graphs: Reliable network computing. In *ICALP (1994)*, vol. 820 of LNCS, Springer-Verlag, pp. 508–515. (preliminary version of [28]).
- [28] NIKOLETSEAS, S., PALEM, K., SPIRAKIS, P. G., AND YUNG, M. Connectivity properties in random regular graphs with edge faults. *International Journal of Foundations of Computer Science (IJFCS)* 11, 2 (2000), 247–262.

- [29] PELEG, D. Proximity-preserving labeling schemes. *Journal of Graph Theory* 33 (2000), 167–176.
- [30] PELEG, D., AND UPFAL, E. A trade-off between space and efficiency for routing tables. *Journal of the ACM* 36, 3 (1989), 510–530.
- [31] ROBERTSON, N., AND SEYMOUR, P. D. Graph minors. III. planar tree-width. *Journal of Combinatorial Theory* 36 (1984), 49–64.
- [32] RUŽIČKA, P. On efficiency of path systems induced by routing and communication schemes. *Computing and Informatics* 20 (2001), 181–205.
- [33] SANTORO, N., AND KHATIB, R. Labelling and implicit routing in networks. *The Computer Journal* 28, 1 (1985), 5–8.
- [34] TARJAN, R. E. Depth-first search and linear graph algorithms. *SIAM Journal on Computing* 1, 2 (1972), 146–160.
- [35] THORUP, M., AND ZWICK, U. Compact routing schemes. In *SPAA (2001)*, ACM PRESS, pp. 1–10.
- [36] TSE, S. S. H., AND LAU, F. C. M. On the space requirement of interval routing. *IEEE Transactions on Computers* 48, 7 (1999), 752–757.
- [37] VAN LEEUWEN, J., AND TAN, R. B. Interval routing. *The Computer Journal* 30, 4 (1987), 298–307.

Cyril Gavoille is with LaBRI, Université Bordeaux I, 351, cours de la Libération, 33405 Talence cedex, France. E-mail: gavoille@labri.fr

Martin Nežez is with the Department of Computer Science and Engineering, Slovak University of Technology, Ilkovičova 3, 812 19 Bratislava. E-mail: nehez@elf.stuba.sk