On the Complexity of Greedy Routing in Ring-based Peer-to-peer Networks*

[Extended Abstract]

George Giakkoupis
ggiak@cs.utoronto.ca

Vassos Hadzilacos
vassos@cs.utoronto.ca

Department of Computer Science
University of Toronto
Toronto, Canada

ABSTRACT

We investigate the complexity of greedy routing in *uniform ring-based random graphs*, a general model that captures many topologies that have been proposed for peer-to-peer and social networks. In this model the nodes form a ring; for each node u we independently draw the set of distances along the ring from u to its "long-range contacts" from a fixed distribution P (the same for all u), and connect u to the corresponding nodes as well as its ring successor. We prove that, for *any* distribution P, in a graph with n nodes and an expected number of ℓ long-range contacts per node constructed in this fashion, the expected number of steps for greedy routing is $\Omega((\log^2 n)/\ell a^{\log^* n})$, for some constant $a > 1$. This improves an earlier lower bound of $\Omega((\log^2 n)/\ell \log\log n)$ by Aspnes *et al.* and is very close to the upper bound of $O((\log^2 n)/\ell)$ achieved by greedy routing in Kleinberg's (one-dimensional) "small-world" networks, a particular instance of uniform ring-based random graphs.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms and Problems**]: Routing and layout; E.1 [**Data Structures**]: Graphs and networks; G.2.2 [**Graph Theory**]: Graph algorithms, Network problems; C.2.2 [**Network Protocols**]: Routing protocols

General Terms

Algorithms, Performance, Theory

Keywords

peer-to-peer, small-worlds, random graphs, greedy routing, lower bound, coupling

*Research supported in part by the Natural Sciences and Engineering Council of Canada.

PODC'07, August 12–15, 2007, Portland, Oregon, USA.

1. INTRODUCTION

Let $\mathbf{G}_n$ be the set of all directed graphs on the set of nodes $[0..n-1]$,[1] where each node u is connected to its *successor* node $(u+1) \bmod n$, and, optionally, to a set of *long-range contacts*. So, every graph in $\mathbf{G}_n$ contains the directed ring $0 \to 1 \to \cdots \to (n-1) \to 0$ as a subgraph. Consider now the following simple model for constructing random graphs in $\mathbf{G}_n$. Fix a probability distribution P on the subsets of $[2..n-1]$, i.e., the set of possible distances *along the ring* from a node to the other nodes, excluding the node's successor. For each node u, independently choose a sample Δ_u from P, and let the long-range contacts of u be the nodes whose distance from u is in Δ_u, i.e., the nodes $(u+d) \bmod n$, for all $d \in \Delta_u$. We call the resulting random graph a *uniform ring-based random graph*, or *uniform graph*, for short.

Uniform graphs capture a wide range of graph topologies proposed for peer-to-peer networks or used to model social networks. Depending on the choice of the distribution P, the resulting graph may be deterministic, such as the Chord ring; or randomized, such as Kleinberg's (one-dimensional) "small-world" networks and randomized versions of Chord.

A natural scheme for routing in networks modeled after uniform graphs is *greedy routing*: given a destination node t, a node forwards a message for t to its neighbor u (successor or long-range contact) that minimizes the distance along the ring to t, i.e., it minimizes $(t-u) \bmod n$.[2] In this paper we investigate the complexity of the above routing scheme in uniform graphs. More precisely, we derive a lower bound on the expected number of steps required to route a message from a randomly selected source to a random destination, in terms of the number of nodes n and the *expected* number ℓ of long-range contacts per node. Note that we do not impose any restrictions on the structure of the distribution P. In particular, the nodes need *not* all have the same number of long-range contacts; also, the long-range contacts of a node are *not* necessarily selected independently of each other.

In his seminal work on routing in social networks [8], Kleinberg described a simple instance of the uniform graphs family where each node u has the same number $\ell \leq \log n$ of long-range contacts, and the distance (along the ring)

[1] For $i, j \in \mathbb{Z}$, we denote by $[i..j]$ the set $\{k \in \mathbb{Z} : i \leq k \leq j\}$.
[2] An alternative version (bidirectional greedy routing) is to forward the message to the neighbor u that minimize the "absolute distance" $|t-u|$ to t. We discuss this variation later.

from u to each of them is drawn independently from the harmonic distribution. The expected number of steps for greedy routing in this construction is $O((\log^2 n)/\ell)$. (This complexity was later shown to be tight [3].) No construction that is modeled after uniform graphs and achieves better than $\Theta((\log^2 n)/\ell)$ has been proposed yet. On the other hand, Aspnes et al. [2] proved that, regardless of the actual distribution P used to determine the long-range contacts, the expected number of steps for greedy routing is $\Omega((\log^2 n)/\ell \log\log n)$.

In this paper we reduce the gap between these two results by improving the lower bound to $\Omega((\log^2 n)/\ell a^{\log^* n})$, for some constant $a > 1$. Note that the quantity $a^{\log^* n}$ grows slower than any constant number of iterative applications of log to n; it is "practically" a constant. The proof of this result is inspired, to some degree, by Kleinberg's analysis. It proceeds by deriving a recursive formula that bounds the expected number of steps for greedy routing on any uniform graph of a given size in terms of that on an exponentially smaller uniform graph. Our analysis suggests general structural properties of the optimal uniform graph that are similar to those observed by Kleinberg for the class of graphs he studied. We conjecture that the lower bound can be further improved to $\Omega((\log^2 n)/\ell)$ — i.e., that Kleinberg's construction is in fact (asymptotically) optimal for greedy routing on uniform graphs.

Discussion and related work. Our results are limited by two assumptions: First, we focus on uniform ring-based random graphs; second, we focus on greedy routing. We now explain the benefits of these assumptions.

Despite its simplicity, the model of uniform graphs can describe a wide range of systems, by suitable choice of the probability distribution P that determines the long-range contacts. For example, we can describe deterministic constructions (where one subset of $[2..n-1]$ has probability one, and all others have probability zero), as well as probabilistic ones. Likewise, we can describe both homogeneous systems, where all nodes have the same number of long-range contacts, and heterogeneous ones where the number of long-range contacts can vary (a little or a lot) from node to node. The uniformity property, i.e., that each node uses independently the same distribution to determine its long-range contacts, is also beneficial. It implies that a node's position in the ring doesn't influence the distances to its long-range contacts. Since nodes are effectively equivalent, it is easier for them to share load and harder for an adversary to disrupt the system by attacking critical nodes.

The advantages of greedy routing are well-known, and reflected by its popularity [1]: Routing decisions are made locally and independently in each node. They are also independent of the routing path up to the current node, so, messages need not store routing information other than the destination node. As a result of these two properties, greedy routing is easy to implement. Also, it is inherently fault-tolerant since as long as each node has *some* edge towards the destination, the message will reach it. Last, it has good locality behavior in that every step reduces the distance to the destination. *Bidirectional* greedy routing is a variation of the (unidirectional) version that we consider in this paper, where a node forwards a message for node t to its neighbor u that minimizes the "absolute distance" $|t-u|$ to t.[3]

In view of the advantages of the two assumptions underlying our analysis, it is not surprising that the designs of many peer-to-peer systems fall within the purview of these assumptions. For both unidirectional and bidirectional greedy routing, there are deterministic constructions with $\ell = \Theta(\log n)$ (Chord [15, 5]), and probabilistic ones with ℓ ranging from 1 to $\Theta(\log n)$ (Kleinberg's small-world networks [8], Symphony [12], Randomized-Chord [6, 17]). In all of these constructions, the expected number of steps for routing is $\Theta((\log^2 n)/\ell)$.

On the lower bound side, the following facts are known about the number of steps for *unidirectional* routing in uniform graphs. Xu has shown that for any *deterministic* construction with $\ell = \Theta(\log n)$, the number of steps is $\Omega(\log n)$ *in the worst-case* [16]. (I.e., there is a source–destination pair of nodes in the given graph for which unidirectional greedy routing requires $\Omega(\log n)$ steps.) Aspnes *et al.* have shown that, for *any* distribution P, the expected number of steps (for a randomly selected source–destination pair of nodes) is $\Omega((\log^2 n)/\ell \log\log n)$ [2]. We improve this bound to $\Omega((\log^2 n)/\ell a^{\log^* n})$, for some constant $a > 1$. For *bidirectional* routing in uniform graphs, Aspnes *et al.* showed that under some assumptions on P, the expected number of steps is $\Omega((\log^2 n)/\ell^2 \log\log n)$ [2]. Also, Flammini *et al.* have shown that in the special case where each node has exactly one long-range contact and certain assumptions on its distribution apply, the expected number of steps for the *worst-case* source–destination pair is $\Omega(\log^2 n)$ [4].

The shortest tree of degree ℓ spanning n nodes has height $\Theta(\log n/\log \ell)$. Thus, this bound represents the optimal expected number of steps for routing in an n-node network of degree ℓ. Indeed, designs that achieve this bound have been proposed. Note that this bound is (asymptotically) better than the lower bound for greedy routing in uniform graphs. In particular, the lower bound of Aspnes *et al.* implies that, for $\ell = o(\log n)$, greedy routing in uniform graphs is suboptimal; that bound, however, is not strong enough to show that this is also true in the (important) case where $\ell = \Theta(\log n)$. Our improved lower bound shows that greedy routing in uniform graphs is suboptimal even in that case.

In view of the lower bound showing that greedy routing in uniform graphs is suboptimal, designs that achieve optimal routing must abandon at least one of the assumptions underlying that bound: Either they must be based on constructions that are not uniform ring-based random graphs, or they must use non-greedy routing — or both. As we argued earlier, these assumptions have considerable advantages, and so the gain of more efficient routing has to be weighed against the loss of these advantages.

We now review some of the designs that achieve more efficient routing than is possible with greedy routing in uniform graphs. Papillon [1] is an example of a network that achieves optimal routing by abandoning only the first assumption. It uses greedy routing but the underlying construction is a *non-uniform* ring-based graph (a ring-embedded butterfly-like network). A similar construction is also described in [4]. The so-called "neighbors-of-neighbors" approach of Manku *et al.* [13] is an example of a network that achieves better routing performance by abandoning only the second assumption. It uses a non-greedy algorithm (where the routing

[3]To ensure that in uniform graphs bidirectional greedy routing can always reach the destination, we slightly modify our model by requiring that, in addition to its successor, each node u is always connected to its *predecessor* $(u-1) \bmod n$.

decision at each node is based not only on the node's neighbors, but also on *their* neighbors), but the network itself is a uniform ring-based random graph. (In fact, the underlying construction is Kleinberg's small-world network). Finally, several deterministic and randomized constructions have been proposed that achieve optimal routing by abandoning both assumptions: They use non-greedy algorithms in networks that are not uniform ring-based random graphs, such as the de Bruijn graph, randomized versions of it, as well as randomized versions of the butterfly network that have been proposed for peer-to-peer networks [10, 7, 11, 14].

A number of recent papers study the problem of constructing random graphs where greedy routing is efficient (i.e., it requires a number of steps at most polylogarithmic in n), by "augmenting" base structures other than the ring (or the grid). In particular, each node of the base graph is augmented by a long-range contact selected independently from some distribution over the remaining nodes (possibly, a different distribution for each node). In the greedy routing scheme they consider, a node forwards a message to its neighbor that has the shortest path to the destination as measured in the base graph (not in the augmented one). For a recent survey of this work see [9].

2. RIGOROUS STATEMENT OF RESULT

For every graph $G \in \mathbf{G}_n$ and pair of nodes u, v of G, we denote by $U(G, u, v)$ the length of the (unidirectional) greedy routing path from u to v; we call it the *routing cost* from u to v. The routing cost from u to the "average destination", where each node is equally likely to be the destination, is denoted $\bar{U}_u(G)$; i.e., $\bar{U}_u(G) = (1/n)\sum_{v=0}^{n-1} U(G, u, v)$.

Let P be a probability distribution on the subsets of $[2..n-1]$. We denote by $\mathbb{G}(P)$ the uniform graph model where the distances to the long-range contacts of a nodes are distributed according to P. The *cardinality* ℓ of P is the expected number of long-range contacts per node in $\mathbb{G}(P)$; i.e., $\ell = \sum_{\Delta \subseteq [2..n-1]} |\Delta| P(\Delta)$. We say that P is an (n, ℓ)-*distribution* if P is a distribution on the subsets of $[2..n-1]$ of cardinality ℓ. The expected routing cost in $\mathbb{G}(P)$ (from a fixed source) to the average destination is denoted $T(P)$; i.e., $T(P) = \mathbb{E}[\bar{U}_u(G)]$, where G is a random graph in $\mathbb{G}(P)$. (Note that because of the way the long-range contacts of G's nodes are selected, $\mathbb{E}[\bar{U}_u(G)]$ is the same for all u.) The optimal $T(P)$ over all possible (n, ℓ)-distributions P is denoted $T(n, \ell)$. (More formally, $T(n, \ell)$ is the *infimum* of $T(P)$ over all P.)

We can now state our main result. Here we assume ℓ is a non-decreasing function of n, perhaps a constant.[4]

THEOREM 1. *If* $\ell = \Omega(1)$ *then* $T(n, \ell) = \Omega((\log^2 n)/\ell a^{\log^* n})$, *where* a *is a constant* > 1.

The corresponding upper bound (which follows mostly from previously known results) is:

THEOREM 2. *If* $\ell = O(\log n)$ *then* $T(n, \ell) = O((\log^2 n)/\ell)$.

3. ANALYSIS

In Section 3.1, we state four auxiliary lemmata and use them to derive Theorems 1 and 2. In Section 3.2, we introduce routing trees, a concept useful for our analysis. Section 3.3 contains the proofs of the first three auxiliary lemmata from Section 3.1. In Section 3.4 we provide some additional facts about routing trees, which we use in Section 3.5 to prove the fourth auxiliary lemma.

3.1 Statement of auxiliary lemmata and derivation of Theorems 1 & 2

We begin with three lemmata that allow us to bound $T(n, \ell)$ in terms of $T(n', \ell')$, for $n' \neq n$ or $\ell' \neq \ell$. The first states the intuitive result that $T(n, \ell)$ is a non-increasing function of ℓ.

LEMMA 1. *If* $\ell < \ell'$ *then* $T(n, \ell) \geq T(n, \ell')$.

The next lemma says what happens to $T(n, \ell)$ for fixed ℓ, as n increases. One might expect $T(n, \ell)$ to be a non-decreasing function of n. This is not necessarily the case, since some "convenient" values of n (say, powers of 2) may result in smaller T than smaller values of n. We show the following weaker result, which, however, suffices for our analysis.

LEMMA 2. *If* $n > n'$ *then* $T(n, \ell) \geq \frac{n'}{n} T(n', \ell)$.

The third lemma is more interesting than the previous two. Lemma 1 shows that, for fixed n, as ℓ increases $T(n, \ell)$ decreases; Lemma 3 says it does not decrease too much.

LEMMA 3. *If* $\ell > \ell'$ *then* $T(n, \ell) \geq \frac{\ell'}{\ell} T(n, \ell')$.

The main part of our analysis is the proof of the following lemma, which gives a lower bound for T when $\ell = 1$. We use $T(n)$ as a shorthand for $T(n, 1)$.

LEMMA 4. $T(n) = \Omega((\log^2 n)/a^{\log^* n})$, *for some constant* $a > 1$.

We can now derive Theorems 1 and 2 as follows.

PROOF OF THEOREM 1. Since ℓ is a non-decreasing function of n, for all sufficiently large values of n it is $\ell < 1$, or $\ell > 1$, or $\ell = 1$. If $\ell < 1$ the bound for $T(n, \ell)$ follows from Lemmata 1 (for $\ell' = 1$) and 4, and the fact that ℓ is larger than some positive constant (since $\ell = \Omega(1)$). If $\ell > 1$ the bound follows from Lemmata 3 (for $\ell' = 1$) and 4. The case $\ell = 1$ is handled in Lemma 4. □

PROOF OF THEOREM 2. Kleinberg showed that $T(n, \ell) = O((\log^2 n)/\ell)$ for $\ell = 1$ [8]. Using similar techniques, Aspnes *et al.* showed that the same bound holds if ℓ takes integer values in the range $1 \leq \ell \leq \log n$ [2]. By Lemma 1, $T(n, \ell) \leq T(n, \lfloor \min\{\ell, \log n\} \rfloor)$, for $\ell \geq 1$. This, together with Aspnes *et al.*'s result, yields the desired bound for all ℓ that take real values such that $1 \leq \ell = O(\log n)$. For $\ell < 1$ the bound follows from Kleinberg's result and the fact that $T(n, \ell) \leq T(n)/\ell$, which follows from Lemma 3. □

3.2 Routing trees

Let $G \in \mathbf{G}_n$ and u be a node of G. Consider the subgraph R of G that consists of the (unidirectional greedy) routing paths from u to all the other nodes. We call R the *routing tree of* G *rooted at* u. An example is illustrated in Figures 1(a)–(c). It is easy to verify that R is indeed a tree. Also, a pre-order walk of R visits the nodes in the order $u, u+1, \ldots, n-1, 0, 1, \ldots, u-1$.[5] Note that the depth of a

[4]Technically, the following weaker condition on ℓ suffices: $\ell = \Theta(g)$, where g is a non-decreasing function of n.

[5]We assume that the children of each internal node of R are ordered from left to right in increasing distance from u.

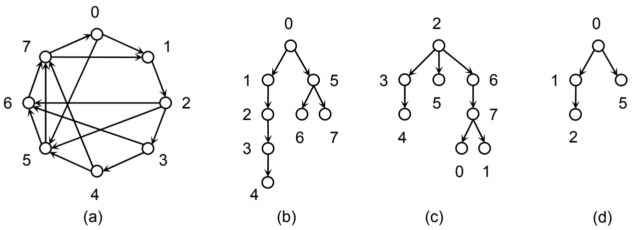

Figure 1: (a) A sample $G \in \mathbf{G}_8$; (b) the routing tree of G (rooted at 0); (c) the routing tree of G rooted at 2; (d) the routing tree of $(\{5\}, \emptyset)$, i.e., of the delta-sets in G of nodes 0 and 1.

node v in R is the routing cost $U(G, u, v)$, and, thus, $\bar{U}_u(G)$ is equal to the average node depth in R.

We use the following conventions. We refer to the routing tree of G rooted at 0 as "the routing tree of G" (omitting the reference to root 0), and to the routing path/cost from 0 to u as "the routing path/cost to u" (omitting the reference to source 0). Likewise, we omit the subscript u from $\bar{U}_u(G)$ when $u = 0$.

We call the set of the distances along the ring from u to its long-range contacts the *delta-set* of u in G. E.g., for the graph in Figure 1(a), the delta-sets of nodes 1 and 2 are $\emptyset$ and $\{3, 4\}$, respectively.

Suppose now we have only partial knowledge of G; in particular, suppose we just know its size n, and the delta-sets $\Delta_0, \Delta_1, \dots, \Delta_{t-1}$ of the first $t \le n$ nodes $0, 1, \dots, t-1$. What can we infer about the routing tree R of G (rooted at 0)? It is easy to see that one can construct the subgraph R^t of R induced by nodes $0, 1, \dots, t-1$ and their children in R. (For $t = 0$, i.e., when we are given no delta-sets, R^t consists only of node 0.) R^t is itself a tree; we call it the *routing tree of* $(\Delta_0, \Delta_1, \dots, \Delta_{t-1})$. An example of such a tree is depicted in Figure 1(d). Each node of G that is not in R^t is in a subtree of R rooted at some leaf $u \ge t$ of R^t. In particular, the subtree of R rooted at such a node u contains the nodes $u, u+1, \dots, u'-1$, where u' is the smallest node in R^t that is larger than u, or $u' = n$ if no such node exists. The size $u' - u$ of this subtree is called the *span of* u *in* R^t. For example, in the routing tree of Figure 1(d) the span of both nodes 2 and 5 is 3.

3.3 Proofs of Lemmata 1–3

All three proofs employ the same technique, based on the coupling method: For any (n, ℓ)-distribution P, we construct a coupling (G, G') of $\mathbb{G}(P)$ and $\mathbb{G}(P')$ for some (n', ℓ')-distribution P' (for suitable n', ℓ', depending on the lemma). In other words, we define random graphs G and G' on a *common* probability space, such that G has the same distribution as $\mathbb{G}(P)$, and G' the same as $\mathbb{G}(P')$. For this pair of random graphs, we show that $\mathbb{E}[\bar{U}(G)] \ge \alpha\, \mathbb{E}[\bar{U}(G')]$ (for a suitable α, depending on the lemma). In the actual construction of G and G', we build G first, according to the definition of $\mathbb{G}(P)$. Then, based on G (and, possibly, on additional random choices) we construct G' such that the following two objectives are satisfied:

Objective 1. G' has the same distribution as $\mathbb{G}(P')$.

Objective 2. $\mathbb{E}[\bar{U}(G') \mid G] \le (1/\alpha)\bar{U}(G)$.

Note that Objective 2 implies $\mathbb{E}[\bar{U}(G)] \ge \alpha\, \mathbb{E}[\bar{U}(G')]$. From this and Objective 1 we obtain $T(P) \ge \alpha T(P')$, and since we assumed an arbitrary P, $T(n, \ell) \ge \alpha T(n', \ell')$. The construction of G' is simple for Lemmata 1 and 2; it is less intuitive for Lemma 3. Below we denote by Δ_u, for $u < n$, the delta-set of node u in G, and by Δ'_u, for $u < n'$, the delta-set of node u in G'.

Proof of Lemma 1. For this lemma $n' = n$ and $\alpha = 1$. Let P' be a distribution obtained from P by reducing the probability associated with (some of) the *proper* subsets of $[2..n-1]$, and correspondingly increasing the probability of the entire set $[2..n-1]$, such that the cardinality of the resulting distribution is ℓ'. It is easy to verify that we can always construct such a P' (for any P and $\ell' > \ell$). G' is defined as follows. For every node u such that $\Delta_u = [2..n-1]$ (i.e., u has outgoing edges to all other nodes in G), $\Delta'_u = \Delta_u$; for each of the remaining u, we independently set $\Delta'_u = \Delta_u$ with probability $P'(\Delta_u)/P(\Delta_u)$, or $\Delta'_u = [2..n-1]$ with probability $1 - P'(\Delta_u)/P(\Delta_u)$. Obviously, for all u, the routing path to u in G' is a prefix of the corresponding routing path in G; thus, $U(G', 0, u) \le U(G, 0, u)$. From this it follows that $\bar{U}(G') \le \bar{U}(G)$ (which implies Objective 2). It is straightforward to verify that Objective 1 holds, as well. □

Proof of Lemma 2. We construct G' by letting $\Delta'_u = \{x \in \Delta_u : x < n'\}$. Clearly, the routing path to $u < n'$ in G' is identical to the corresponding routing path in G, thus, $U(G', 0, u) = U(G, 0, u)$. From this, $\bar{U}(G') = (1/n') \cdot \sum_{u<n'} U(G, 0, u) \le (1/n') \sum_{u<n} U(G, 0, u) = (n/n')\bar{U}(G)$ (which implies Objective 2, for $\alpha = n'/n$). It is also easy to verify that Objective 1 holds, for a P' of cardinality $\ell' \le \ell$. Therefore, $T(n, \ell) \ge (n'/n)T(n', \ell')$, and since $T(n', \ell') \ge T(n', \ell)$ (by Lemma 1), $T(n, \ell) \ge (n'/n)T(n', \ell)$. □

Proof of Lemma 3. For this lemma $n' = n$ and $\alpha = \ell'/\ell$. We first describe the construction of G' informally. In this construction, we associate with each node u of G' a node C_u of G (this association is not necessarily one-to-one). G' is constructed inductively by considering each node $u = 0, 1, \dots, n-1$ in turn. Let R and R' be the routing trees of G and G', respectively. We initialize the inductive construction by setting $C_0 = 0$ — i.e., associating with the root of R' the root of R. For the node u of G' under consideration, we define its delta-set Δ'_u and, simultaneously, define the association of u's children in R' to nodes in R. Specifically, we choose $\Delta'_u = \emptyset$ with probability $1 - \alpha$, and $\Delta'_u = \Delta_{C_u}$ with probability α. So, in the first case u has no long-range contacts in G' and (if u is not a leaf of R') its only child in R' is $u + 1$. In this case, we associate with $u + 1$ the node in G to which u is already associated, i.e., $C_{u+1} = C_u$. In the second case, u has long-range contacts in G' at the same distances as the corresponding node C_u does in G and this defines the set of u's children in R'. In this case, we associate with each child $u + \delta$ of u in R' the corresponding child of C_u in R, i.e., $C_{u+\delta} = C_u + \delta$.

The construction of G' is formalized by the following inductive definitions of Δ'_u and C_u. Let $B_0, B_1, \dots, B_{n-1}$ be independent identically-distributed binary random variables such that $\mathbb{Pr}[B_u = 1] = \alpha$. ($B_0, B_1, \dots, B_{n-1}$ are also independent of G.) For each $u \in [0..n-1]$,

$$\Delta'_u = \begin{cases} \emptyset, & \text{if } B_u = 0 \\ \Delta_{C_u}, & \text{if } B_u = 1 \end{cases}$$

where:

- $C_0 = 0$, and if $u > 0$

$$C_u = \begin{cases} C_{F_u}, & \text{if } B_{F_u} = 0 \\ C_{F_u} + (u - F_u), & \text{if } B_{F_u} = 1 \end{cases}$$

- F_u is u's parent in the routing tree of $(\Delta'_0, \ldots, \Delta'_{u-1})$.

We first establish that G' is well defined. This is immediate from the following fact. (The proof is an easy induction and is omitted).

FACT 3.1. *For all $u \in [0..n-1]$, $0 \le C_u \le u$.*

The next fact gives more insight into the construction of G'. (The proof, by induction, is omitted). Let S_u denote the size of the subtree of R rooted at u, and H_u the size of the subtree of R' rooted at u. Recall that the sets of nodes of the above subtrees are $[u..u + S_u - 1]$ and $[u..u + H_u - 1]$, respectively.

FACT 3.2. *For all $u \in [0..n-1]$,*

(a) $H_u \le S_{C_u}$

(b) if $u < u' < u + H_u$ then $C_u \le C_{u'} < C_u + S_{C_u}$; in particular, if $B_u = 1$ then $C_{u'} > C_u$;
if $u + H_u \le u' < n$ then $C_{u'} \ge C_u + S_{C_u}$

Part (b) says that if $u' > u$ then $C_{u'} \ge C_u$ (where the inequality is strict if $B_u = 1$); also, $C_{u'}$ is a descendant of C_u in R iff u' is a descendant of u in R'.

Next, we show that Objective 1 holds; i.e., that G' has the desired distribution. For this, it is convenient to think of the construction of G' as a random process consisting of n steps, 0 up to $n-1$, where step u determines the value of Δ'_u. We assume that the value of each B_v and Δ_v is generated right before it is about to be used — not earlier. Clearly, B_u is generated in step u. Let $\mathcal{D}_u$ be the subset of G's nodes consisting of all v for which Δ_v is generated in steps 0 up to $u - 1$. Observe that $\mathcal{D}_u = \{C_v : v < u,\ B_v = 1\}$, so, Fact 3.2(b) implies that $C_u \notin \mathcal{D}_u$; i.e., Δ_{C_u} is not generated before step u. Therefore, in each step u: we first determine B_u (independently of past choices); if $B_u = 0$ we set $\Delta'_u = \emptyset$; otherwise, we determine Δ_{C_u} (again independently of past decisions) and let $\Delta'_u = \Delta_{C_u}$. Consequently, $\Delta'_0, \Delta'_1, \ldots, \Delta'_{n-1}$ are distributed independently and identically, and their common distribution P' is

$$P'(\Delta) = \begin{cases} \alpha P(\Delta), & \text{if } \Delta \ne \emptyset \\ (1-\alpha) + \alpha P(\emptyset), & \text{if } \Delta = \emptyset \end{cases}$$

The cardinality of P' is $\sum_\Delta |\Delta| P'(\Delta) = \sum_{\Delta \ne \emptyset} |\Delta| P'(\Delta) = \alpha \sum_{\Delta \ne \emptyset} |\Delta| P(\Delta) = \alpha\ell = \ell'$, as desired.

It remains to show Objective 2. For $t \le n$, let R^t be the routing tree of $(\Delta'_0, \Delta'_1, \ldots, \Delta'_{t-1})$, and V^t be the set of nodes of R^t. We now define a sequence $X_0, X_1, \ldots, X_n$ of progressively more accurate estimates of $n\bar{U}(G')$, where estimate X_t is based on R^t. Specifically, for each $t \in [0..n]$,

$$X_t = \sum_{u \in [0..t-1]} U(G', 0, u) + \sum_{u \in V^t - [0..t-1]} \Big(H_u \cdot U(G', 0, u) + (1/\alpha) \cdot \sum_{v \in [C_u..C_u + H_u - 1]} U(G, C_u, v) \Big)$$

(Recall that $[0..t-1] \subseteq V^t$, and for all $u \in V^t$, $U(G', 0, u)$, H_u, and C_u are functions of R^t; in particular, H_u is equal to the span of u in R^t, for all $u \in V^t - [0..t-1]$.) In the definition of X_t, the first sum accounts for the routing costs to the first t nodes of G'; the second is an estimate of the routing costs to the remaining nodes of G', i.e., to the nodes in the subtrees of R' rooted at all $u \ge t$ that are leaves of R^t. In particular, inside this second sum, the first term accounts for the routing costs up to u for all (the H_u) nodes in R''s subtree rooted at u; the second term is an estimate of the remaining routing costs from u to these nodes. This estimate is proportional to the sum of the routing costs in G from C_u to nodes $C_u, C_u + 1, \ldots, C_u + H_u - 1$, which, by Fact 3.2(a), are all in the subtree of R rooted at C_u. Finally, note that $X_0 = (n/\alpha)\bar{U}(G)$ and $X_n = n\bar{U}(G')$.

We now show that $\mathbb{E}[X_{t+1} \mid G, R^t] \le X_t$ (i.e., that conditioned on G, the sequence of X_t is a super-martingale). Let $Z_{t+1} = X_{t+1} - X_t$. It is not hard to verify that

$$Z_{t+1} = \begin{cases} H_t - 1 - \frac{1}{\alpha} U(G, C_t, C_t + H_t - 1), & \text{if } B_t = 0 \\ -(\frac{1}{\alpha} - 1)(H_t - 1), & \text{if } B_t = 1 \end{cases}$$

Let A_t stand for "G, R^t". Then,

$$\begin{aligned} \mathbb{E}[Z_{t+1} \mid A_t] &= \mathbb{E}[Z_{t+1} \mid A_t, B_t = 0] \cdot \mathbb{Pr}[B_t = 0 \mid A_t] \\ &\quad + \mathbb{E}[Z_{t+1} \mid A_t, B_t = 1] \cdot \mathbb{Pr}[B_t = 1 \mid A_t] \\ &= \big(H_t - 1 - \tfrac{1}{\alpha} U(G, C_t, C_t + H_t - 1)\big)(1 - \alpha) \\ &\quad - (\tfrac{1}{\alpha} - 1)(H_t - 1)\alpha \\ &= -(\tfrac{1}{\alpha} - 1) \cdot U(G, C_t, C_t + H_t - 1) \le 0 \end{aligned}$$

So, $\mathbb{E}[Z_{t+1} \mid G, R^t] \le 0$, or, equivalently, $\mathbb{E}[X_{t+1} \mid G, R^t] \le X_t$. From this it follows that $\mathbb{E}[X_n \mid G] \le X_0$; substituting the values of X_0 and X_n and dividing by n we obtain Objective 2. □

3.4 More on routing trees

We now introduce some additional terminology related to routing trees, which we use in the proof of Lemma 4. Let R be the routing tree of $G \in \mathbf{G}_n$. Let also R_u, for $u < n$, be the subtree of R rooted at node u, and s_u be the size of R_u. A node v of G is an *r-descendant* of u, for some $r \ge 1$, if v is in R_u and $v - u < r$; so, the r-descendants of u are the nodes in $\langle u, \min\{r, s_u\}\rangle$, where by $\langle i, k\rangle$, for $i, k \in \mathbb{Z}$, we denote the set $[i..i + k - 1]$. If v's parent in R is an r-descendant of u, but v itself is not then v is called an *r-successor* of u. So, if $s_u \le r$ all the nodes of R_u are r-descendants of u — u has no r-ancestors; if $s_u > r$ the r-descendants and r-successors of u are as follows: Let $\mathbf{x} = x_0, x_1, \ldots, x_k$ be the path in R_u from u to node $u + r - 1$ (that is, the largest r-descendant of u). The r-descendants of u are the nodes along $\mathbf{x}$, plus the nodes of the subtrees rooted at the children of nodes $x_0, x_1, \ldots, x_{k-1}$ that lie to the *left* of $\mathbf{x}$ on the plane; the r-successors of u are the children of nodes $x_0, x_1, \ldots, x_{k-1}$ that lie to the *right* of $\mathbf{x}$, plus all the children of x_k. We call $\mathbf{x}$ the *r-border* of u. Examples of these definitions are illustrated in Figure 2(a). Note that the r-successors of u form a "frontier" between the r-descendants of u and the other nodes of R_u: the path from u to each of its r-descendants consists only of r-descendants of u; the path to any other node of R_u consists of one or more r-descendants of u followed by exactly one r-successor of u and then zero or more nodes that are neither r-descendants nor r-successors of u.

Consider the set of nodes that consists of node 0, 0's r-successors, the r-successors of them, and so on. The nodes in this set are called the *r-significant* nodes of G. Note that if $z_0 < z_1 < \cdots < z_{\kappa-1}$ are the r-significant nodes of G and $z_\kappa = n$ then the r-descendants of z_k, for $k < \kappa$, are the nodes

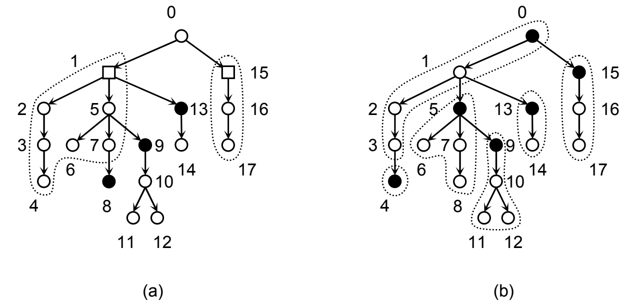

Figure 2: (a) the 7-descendants of 1 are nodes 1–7, and of 15 are nodes 15–17; the 7-successors of 1 are nodes 8, 9, 13; the 7-border of 1 is path 1, 5, 7; (b) the 4-significant nodes are marked with filled circles; the 4-partition of $\langle 0, 12\rangle$ is $\{[0..3], \{4\}, [5..8], [9..11]\}$.

$z_k, z_k+1, \ldots, z_{k+1}-1$. If v is an r-descendant of some r-significant node u, we say that u is the *r-ancestor* of v. Let Π be the partition of set $\langle 0, m\rangle \subseteq \langle 0, n\rangle$ induced by the "have the same r-ancestor" relation. So, for $z_0, z_1, \ldots, z_\kappa$ as above, $\Pi = \{[z_0..z_1-1], [z_1..z_2-1], \ldots, [z_{\kappa'-1}..m-1]\}$, where κ' is such that $z_{\kappa'-1} < m \leq z_{\kappa'}$. We call Π the *r-partition of $\langle 0, m\rangle$ with respect to G*. An example is depicted in Figure 2(b).

From the discussion above, the routing path in G to a node v can be divided into a number of smaller paths, where each smaller path consists of an r-significant node followed by zero or more r-descendants of that node. For example, if the routing tree of G is as in Figure 2(b) and $r = 4$, the routing path to node 12 is divided into the three paths: $0, 1$; 5; and $9, 10, 12$. We can identify the r-significant nodes along the routing path to v in an "on-line" fashion, as follows. Suppose that initially we have no knowledge of G (other than its size n), and in each step we learn the delta-set of the next node in the routing path to v (starting from 0's). Before the first step we mark node 0; whenever a node u is marked we check if $v - u \geq r$; if this condition holds we continue to take steps until we find the next node u' in the path such that $u' - u \geq r$, and we mark u' as well; if it is $v - u < r$ the process ends. The nodes marked are the r-significant nodes contained in the routing path to v. In particular, the last of the nodes marked is the r-ancestor of v. Likewise, we can compute the number m' of the r-descendants in G of any r-significant node u' in the routing path to v, knowing just the delta-sets of the nodes that precede u' in this path. It is easy to verify that $m' = \min\{r, b - u'\}$, where b is the minimum among the right siblings of the nodes that precede u' in the path, or $b = n$ if no such siblings exist.

3.5 Proof of Lemma 4

Let P be any $(n, 1)$-distribution, and G be a random graph in $\mathbb{G}(P)$. Throughout the proof we use U_u as a shorthand for $U(G, 0, u)$. Also by $\bar{U}_m$, for $1 \leq m \leq n$, we denote the average routing cost to the nodes in $\langle 0, m\rangle$; i.e., $\bar{U}_m = (1/m)\sum_{u\in\langle 0,m\rangle} U_u$.

Roughly, the proof proceeds as follows. First, we show that for any positive integers m, r, η such that $\eta \leq r \leq m \leq n$, $\mathbb{E}[\bar{U}_m]$ is bounded by the sum of the two terms:

(i) $\mathbb{E}[U_Z]$, where Z is the r-ancestor of a random node Y in $\langle 0, m\rangle$; and

(ii) $\mathbb{E}[\bar{U}_{m^*}]$, for an "optimal" $m^* \in [\eta, r]$, scaled by the probability q that Z has at least η r-descendants in $\langle 0, m\rangle$.

Term (ii) bounds the expected (remaining) routing cost from Z to Y. We then compute a lower bound for $\mathbb{E}[\bar{U}(G)] = \mathbb{E}[\bar{U}_n]$ by recursively applying the above result. In particular, in each recursive step we take $r \approx m/\log^\beta n$ and $\eta \approx r/\log^{\beta+\gamma} n$, for some constants β, γ. For this choice of r and η we show that each term of type (i) involved in the bound for $\mathbb{E}[\bar{U}(G)]$ can be bounded by $T(n')$, for some n' exponentially smaller than n, scaled by a factor proportional to the "density" of P in the range, roughly, between r and m. (By density of P in some range of distances we mean the expected number of long-range contacts per node at distances in this range.) We also show that the probabilities q involved (because of the terms of type (ii)) are very close to 1. By applying a convexity argument to the resulting bound for $\mathbb{E}[\bar{U}(G)]$ we obtain $\mathbb{E}[\bar{U}(G)] \geq T(n')/a$, and, thus, $T(n) \geq T(n')/a$. Recursive application of the last inequality yields the desired result.

The convexity argument we mentioned above suggests that, optimally, the density of P should be *roughly* the same over intervals of the form $[u..u\log^c n]$, for a constant c. Recall that Kleinberg's optimal small-world network has the property that the density of P is *exactly* the same for all intervals of the form $[u..cu]$.

We now describe the details of the proof. Let Π be the r-partition of $\langle 0, m\rangle$ with respect to G, where $r \leq m \leq n$. Let also $\langle Z, M\rangle$ be a random block of Π, where Z is the r-ancestor of a uniformly random node in $\langle 0, m\rangle$, and M is the number of Z's r-descendants in $\langle 0, m\rangle$. Clearly, $\mathbb{P}\mathrm{r}[\langle Z, M\rangle = \langle z, m'\rangle \mid G] = (m'/m)$ if $\langle z, m'\rangle \in \Pi$, and 0 otherwise. Note that

$$\begin{aligned}\bar{U}_m &= \frac{1}{m}\sum_{\langle z,m'\rangle\in\Pi}\ \sum_{u\in\langle z,m'\rangle} U_u \\ &= \sum_{\langle z,m'\rangle\in\Pi}\Big(\mathbb{P}\mathrm{r}[\langle Z, M\rangle = \langle z, m'\rangle \mid G]\cdot\frac{1}{m'}\sum_{u\in\langle z,m'\rangle} U_u\Big) \\ &= \mathbb{E}\Big[\frac{1}{M}\sum_{u\in\langle Z,M\rangle} U_u \;\Big|\; G\Big]\end{aligned}$$

Taking the expectation over G, and observing that it is $\sum_{u\in\langle Z,M\rangle} U_u = M\cdot U_Z + \sum_{u\in\langle Z,M\rangle} U(G, Z, u)$ we obtain

$$\mathbb{E}[\bar{U}_m] = \mathbb{E}[U_Z] + \mathbb{E}\Big[\frac{1}{M}\sum_{u\in\langle Z,M\rangle} U(G, Z, u)\Big] \qquad (1)$$

We focus now on the rightmost term of (1), namely, on $\mathbb{E}[(1/M)\sum_{u\in\langle Z,M\rangle} U(G, Z, u)]$. From the discussion at the end of Section 3.4, Z and M can be determined without knowing the delta-sets of the nodes in $\langle Z, M\rangle$. Therefore, the knowledge of Z and M does not reveal any information about the delta-sets of the nodes in $\langle Z, M\rangle$. More formally, let Δ_u denote the delta-set of u. Conditioned on any event $E_{z,m'} = \{\langle Z, M\rangle = \langle z, m'\rangle\}$ such that $\mathbb{P}\mathrm{r}[E_{z,m'}] > 0$, $\Delta_z, \Delta_{z+1}, \ldots, \Delta_{z+m'-1}$ are independent and identically distributed according to P. Consequently, for all $u \in \langle z, m'\rangle$, $\mathbb{E}[U(G, Z, u) \mid E_{z,m'}] = \mathbb{E}[U_{u-z}]$. So,

$$\begin{aligned}&\mathbb{E}\Big[\frac{1}{M}\sum_{u\in\langle Z,M\rangle} U(G, Z, u)\Big] \\ &\quad= \sum_{z,m'}\Big(\mathbb{E}\Big[\frac{1}{M}\sum_{u\in\langle Z,M\rangle} U(G, Z, u) \;\Big|\; E_{z,m'}\Big]\cdot\mathbb{P}\mathrm{r}[E_{z,m'}]\Big)\end{aligned}$$

$$= \sum_{z,m'} \Big(\frac{1}{m'} \sum_{u \in \langle 0,m'\rangle} \mathbb{E}[U_u] \cdot \mathbb{Pr}[E_{z,m'}]\Big)$$

$$= \sum_{z,m'} \mathbb{E}[\bar{U}_{m'}] \cdot \mathbb{Pr}[E_{z,m'}] = \sum_{m'=1}^{r} \mathbb{E}[\bar{U}_{m'}] \cdot \mathbb{Pr}[M = m']$$

If we restrict the range of m' in the last sum to $\eta \le m' \le r$, where $1 \le \eta \le r$, and let m^* be the $m' \in [\eta..r]$ of minimum $\mathbb{E}[\bar{U}_{m'}]$, we get $\mathbb{E}[(1/M)\sum_{u \in \langle Z,M\rangle} U(G,Z,u)] \ge \mathbb{E}[\bar{U}_{m^*}] \cdot \mathbb{Pr}[M \ge \eta]$. From this and (1) we have

$$\mathbb{E}[\bar{U}_m] \ge \mathbb{E}[U_Z] + \mathbb{E}[\bar{U}_{m^*}] \cdot \mathbb{Pr}[M \ge \eta] \quad (2)$$

The fact we state next provides lower bounds for $\mathbb{E}[U_Z]$ and $\mathbb{Pr}[M \ge \eta]$. Its proof is described in Sections 3.5.1 and 3.5.2. Let $2 \le \theta < r$. Let also λ be the expected number of long-range contacts a node has at distances between θ and $m-1$; i.e., $\lambda = \mathbb{E}[|\Delta_u \cap [\theta..m-1]|]$. (Note that λ is the *density* of P in $[\theta..m-1]$, as defined at the beginning of the proof.) Suppose also $m \ge 3r$.

FACT 4.1.
(a) $\mathbb{E}[U_Z] \ge c \cdot \min\{r/\theta,\ \frac{1}{\lambda}T(\lfloor \frac{m}{r} \rfloor)\}$,
(b) $\mathbb{Pr}[M \ge \eta] \ge 1 - c'\,\mathbb{E}[U_{r-1}] \cdot (\eta - 1)/r$,
where c, c' are positive constants.

We will now use result (2) and Fact 4.1 to compute a bound for $\mathbb{E}[\bar{U}(G)]$. Let $n_0 = n$, and for each $i \in [0..\tau-1]$, where $\tau = \min\{j : n_j \le \log^{\beta+\gamma} n\}$,

- r_i is the $k \in [\lceil n_i/\log^\beta n\rceil..\lceil 2n_i/\log^\beta n\rceil]$ that minimizes $\mathbb{E}[U_{k-1}]$
- $\eta_i = \lceil n_i/\log^{\beta+\gamma} n\rceil$
- $\theta_i = \lceil n_i/\log^{\beta+\delta} n\rceil$
- n_{i+1} is the $k \in [\eta_i..r_i]$ that minimizes $\mathbb{E}[\bar{U}_k]$
- $\lambda_i = \mathbb{E}[|\Delta_u \cap [\theta_i..n_i - 1]|]$
- Z_i is the r_i-ancestor of a uniformly random node in $\langle 0, n_i\rangle$
- M_i is the number of the r_i-descendants of Z_i in $\langle 0, n_i\rangle$

β, γ, and δ are constants such that $\gamma \ge \beta+3$ and $\beta \ge \delta \ge 2$. (Note that all the above quantities except for the last two are computed directly from P.) By recursively applying (2) we obtain the following bound for $\mathbb{E}[\bar{U}(G)] = \mathbb{E}[\bar{U}_n]$. For all $i' \in [0..\tau]$,

$$\mathbb{E}[\bar{U}(G)] \ge \sum_{i<i'} \Big(\mathbb{E}[U_{Z_i}] \cdot \prod_{j<i} \mathbb{Pr}[M_j \ge \eta_j]\Big) + \mathbb{E}[\bar{U}_{n_{i'}}] \cdot \prod_{j<i'} \mathbb{Pr}[M_j \ge \eta_j] \quad (3)$$

From Fact 4.1(a), $\mathbb{E}[U_{Z_i}] \ge c \cdot \min\{r_i/\theta_i,\ (1/\lambda_i)T(\lfloor n_i/r_i\rfloor)\}$. Since $\frac{1}{2}\log^\beta n - 1 \le n_i/r_i \le \log^\beta n$, Lemma 2 implies that $T(\lfloor n_i/r_i\rfloor) \ge (\frac{1}{2} - \mathrm{o}(1)) \cdot T(\lfloor \frac{1}{2}\log^\beta n - 1\rfloor)$. Also, $r_i/\theta_i \ge (1-\mathrm{o}(1))\log^\delta n \ge (1-\mathrm{o}(1))\log^2 n$, since $\delta \ge 2$. Therefore,

$$\mathbb{E}[U_{Z_i}] \ge \hat{c} \cdot \min\big\{\log^2 n,\ \tfrac{1}{\lambda_i} T\big(\lfloor \tfrac{1}{2}\log^\beta n - 1\rfloor\big)\big\} \quad (4)$$

From Fact 4.1(b), $\mathbb{Pr}[M_i \ge \eta_i] \ge 1 - c'\,\mathbb{E}[U_{r_i-1}](\eta_i - 1)/r_i$. By the definition of r_i, $(\lceil 2n_i/\log^\beta n\rceil - \lceil n_i/\log^\beta n\rceil + 1) \cdot \mathbb{E}[U_{r_i-1}] \le n_i \cdot \mathbb{E}[\bar{U}_{n_i}]$, which implies $\mathbb{E}[U_{r_i-1}] \le \log^\beta n \cdot \mathbb{E}[\bar{U}_{n_i}]$. Also, $(\eta_i - 1)/r_i \le 1/\log^\gamma n \le 1/\log^{\beta+3} n$, since $\gamma \ge \beta+3$. Therefore,

$$\mathbb{Pr}[M_i \ge \eta_i] \ge 1 - c'\,\mathbb{E}[\bar{U}_{n_i}]/\log^3 n \quad (5)$$

Using (3)–(5) we bound $\mathbb{E}[\bar{U}(G)]$ as follows. Assume n is large enough that $\tau > 0$; i.e., $n \ge \hat{n} = \min\{k : k > \log^{\beta+\gamma} k\}$. We distinguish two cases.

CASE 1: $\max\{\mathbb{E}[\bar{U}_{n_i}],\ \mathbb{E}[U_{Z_i}]/\hat{c}\} \ge \log^2 n$, for some $i < \tau$. Let i' be the smallest such i. By (5), for all $i < i'$, $\mathbb{Pr}[M_i \ge \eta_i] \ge 1 - c'/\log n$. If $\mathbb{E}[\bar{U}_{n_{i'}}] \ge \log^2 n$ then, by (3), $\mathbb{E}[\bar{U}(G)] \ge \mathbb{E}[\bar{U}_{n_{i'}}] \cdot \prod_{j<i'} \mathbb{Pr}[M_j \ge \eta_j] \ge \log^2 n(1 - c'/\log n)^{i'}$. Similarly, if $\mathbb{E}[U_{Z_{i'}}] \ge \hat{c}\log^2 n$ then $\mathbb{E}[\bar{U}(G)] \ge \mathbb{E}[U_{Z_{i'}}] \cdot \prod_{j<i'} \mathbb{Pr}[M_j \ge \eta_j] \ge \hat{c}\log^2 n(1 - c'/\log n)^{i'}$. Since $i' < \tau \le \log n/\beta\log\log n + \mathrm{o}(1)$, we have that in either case,

$$\mathbb{E}[\bar{U}(G)] \ge c_1 \log^2 n$$

CASE 2: $\max\{\mathbb{E}[\bar{U}_{n_i}],\ \mathbb{E}[U_{Z_i}]/\hat{c}\} < \log^2 n$, for all $i < \tau$. As in the previous case, $\mathbb{Pr}[M_i \ge \eta_i] \ge 1 - c'/\log n$, for all $i < \tau$. So, by (3) for $i' = \tau$,

$$\mathbb{E}[\bar{U}(G)] \ge \sum_{i<\tau} \Big(\mathbb{E}[U_{Z_i}] \cdot \prod_{j<i} \mathbb{Pr}[M_j \ge \eta_j]\Big) \ge \Big(1 - \frac{c'}{\log n}\Big)^\tau \sum_{i<\tau} \mathbb{E}[U_{Z_i}] \ge \hat{c}' \sum_{i<\tau} \mathbb{E}[U_{Z_i}]$$

Note that (4) and case hypothesis $\mathbb{E}[U_{Z_i}] < \hat{c}\log^2 n$ imply that for all $i < \tau$, $\mathbb{E}[U_{Z_i}] \ge \hat{c}(1/\lambda_i) \cdot T(\lfloor \frac{1}{2}\log^\beta n - 1\rfloor)$. Therefore,

$$\mathbb{E}[\bar{U}(G)] \ge \hat{c}'' T(\lfloor \tfrac{1}{2}\log^\beta n - 1\rfloor) \sum_{i<\tau} 1/\lambda_i$$

Recall that $\lambda_i = \mathbb{E}[|\Delta_u \cap [\theta_i..n_i - 1]|]$, and observe that since $\beta \ge \delta$ each $d < n$ belongs to at most two sets $[\theta_i..n_i - 1]$, for $i < \tau$. So, $\sum_{i<\tau} \lambda_i \le 2\,\mathbb{E}[|\Delta_u|] = 2$, which implies that $\sum_{i<\tau} 1/\lambda_i \ge \tau^2/2$ (because of the convexity of $1/x$). Also, $\tau \ge \log n/(\beta+\gamma)\log\log n$. Therefore,

$$\mathbb{E}[\bar{U}(G)] \ge f(n) \cdot T(\alpha(n))$$

where $f(n) = c_2(\log n/\log\log n)^2$ and $\alpha(n) = \lfloor \frac{1}{2}\log^\beta n - 1\rfloor$. {END OF CASE 2}

Combining CASE 1 and 2, and recalling that the above results hold for *any* $(n,1)$-distribution P, we obtain that for all $n \ge \hat{n}$,

$$T(n) \ge \min\big\{c_1 \log^2 n,\ f(n) \cdot T\big(\alpha(n)\big)\big\} \quad (6)$$

The lemma follows by recursive application of the above inequality. Informally, we recursively apply (6) until either (i) in some recursive step the first argument of $\min\{\cdot,\cdot\}$ is smaller than the second; or (ii) $\kappa = \log^* n - \log^* \hat{n} - 1$ steps have been taken. So, there are $\kappa+1$ possible cases we need to consider: one for each of the κ steps in which condition (i) may be met, and the case where the recursion stops because condition (ii) is met (and (i) is not). The corresponding bounds for $T(n)$ are $\prod_{i<k} f(\alpha^{(i)}(n)) \cdot c_1\log^2(\alpha^{(k)}(n))$, for each $k \in [0..\kappa-1]$, and $\prod_{i<\kappa} f(\alpha^{(i)}(n)) \cdot (2/3)$, where in the last bound we used the trivial fact $T(m) \ge 2/3$, for $m \ge 3$. (By $f^{(j)}(x)$ we denote the function $f(x)$ iteratively applied $j \ge 0$ times to an initial value of x.) We show that all these bounds are $\Omega((\log^2 n)/a^{\log^* n})$, for some constant $a > 1$. Roughly speaking, we do so by establishing that for all $k \in [1..\kappa]$, $\alpha^{(k)}(n) = (\sigma_k \log^{(k)} n)^\beta$, and $f(\alpha^{(k)}(n)) = \sigma'_k(\log^{(k+1)} n/\ \log^{(k+2)} n)^2$, where σ_k and σ'_k lie between positive constants. (The details are omitted.)

3.5.1 *Proof sketch of Fact 4.1(a)*

The proof is similar, in principle, to those of Lemmata 1–3: in the same probability space as G we construct an m'-node uniform graph G', where $m' \approx m/r$, and we bound $\mathbb{E}[U_Z]$ in terms of $\mathbb{E}[\bar{U}(G')]$. Roughly speaking, G' is a "scaled-down" version of the m-node uniform graph $\hat{G}$, where each node k of $\hat{G}$ has delta-set $\Delta_k \cap \langle 0, m\rangle$. ($\Delta_k$ is the delta-set of k in G.) To make routing costs in G' comparable to those in $\hat{G}$ we assign to each node u of G' a positive *weight* W_u, and consider, instead, the "weighted" cost of each routing path in G', that is, the sum of the weights of the node in this path. We show that $\mathbb{E}[U_Z]$ is bounded by the expected weighted routing cost in G' to the average destination, which, in turn, we bound in terms of $\mathbb{E}[\bar{U}(G')]$.

Before we describe how to construct G' and the weights $W_0, W_1, \ldots, W_{m'-1}$ in terms of G, we describe the distribution they will have. We also compute some quantities related to this distribution. The size of G' is $m' = \lfloor m/r' \rfloor$, where $r' = r + \theta$. The pairs $(\Delta'_0, W_0), \ldots, (\Delta'_{m'-1}, W_{m'-1})$, where Δ'_u denotes the delta-set of u in G', are independent and identically distributed. Their common distribution is as follows. Let $D_0, D_1, \ldots$ be a sequence of independent random elements distributed according to P. Let also $J_t = \max(\{1\} \cup D_t \cap \langle 0, m\rangle)$, for $t \geq 0$. Define $\tau_1 = \min\{t : \sum_{j=0}^{t} J_j \geq r\}$ and $\tau_2 = \min\{t : J_t \geq \theta\}$. Finally, let $W = \min\{\tau_1, \tau_2\} + 1$, and $D = \Phi(D_{\min\{\tau_1,\tau_2\}})$, where

$$\Phi(\Delta) = \bigcup_{d \in \Delta \cap [\theta..m-1]} \left\{ \left\lfloor \frac{d}{r'} \right\rfloor + \delta : \delta = -1, 0, 1, 2 \right\} \cap [2..m'-1]$$

Note that if $\tau_1 < \tau_2$ then $D = \Phi(D_{\tau_1}) = \emptyset$, while if $\tau_1 \geq \tau_2$ $D = \Phi(D_{\tau_2}) \neq \emptyset$. The joint distribution of D, W is denoted by $\hat{P}'$, and the marginal of D by P'. $\hat{P}'$ will be the joint distribution of Δ'_u, W_u (and P' the marginal of Δ'_u).

The next fact provides lower bounds for the cardinality ℓ' of P' and the expected value of W. Recall that λ is the expected number of long-range contacts a node in $\mathbb{G}(P)$ has at distances between θ and $m-1$. Let π be the probability a node in $\mathbb{G}(P)$ has at least one such long-range contact; i.e., $\pi = \Pr[\Delta_u \cap [\theta..m-1] \neq \emptyset]$. (Note that $\lambda \geq \pi$.)

FACT 4.1.1.

(a) $\ell' \leq 4\lambda/\pi$

(b) $\mathbb{E}[W \mid D] \geq c_1 \min\{r/\theta, 1/\pi\}$*, for a constant* $c_1 > 0$.

PROOF. For part (a) we have

$$\begin{aligned}\ell' &= \mathbb{E}[|D|] = \mathbb{E}\left[|\Phi(D_{\tau_2})| \mid \tau_1 \geq \tau_2\right] \cdot \Pr[\tau_1 \geq \tau_2] \\ &\leq \mathbb{E}\left[|\Phi(D_{\tau_2})| \mid \tau_1 \geq \tau_2\right] = \mathbb{E}\left[|\Phi(D_t)| \mid J_t \geq \theta\right] \\ &= \mathbb{E}[|\Phi(D_t)|] / \Pr[J_t \geq \theta]\end{aligned}$$

where the last equality holds because $\Phi(D_t) = \emptyset$ if $J_t < \theta$. It is easy to verify that $|\Phi(D_t)| \leq 4|D_t \cap [\theta..m-1]|$; so, $\mathbb{E}[|\Phi(D_t)|] \leq 4\,\mathbb{E}[|D_t \cap [\theta..m-1]|] = 4\lambda$. Also, $\Pr[J_t \geq \theta] = \Pr[D_t \cap [\theta..m-1] \neq \emptyset] = \pi$. Therefore, $\ell' \leq 4\lambda/\pi$.
For part (b), for the case where $D = \emptyset$ we have

$$\mathbb{E}[W \mid D = \emptyset] = \mathbb{E}[\tau_1 + 1 \mid \tau_1 < \tau_2] \geq r/\theta + 1$$

The first relation holds because $D = \emptyset$ iff $\tau_1 < \tau_2$; the second because $\tau_1 < \tau_2$ implies $J_t < \theta$ for all $t \leq \tau_1$, and, thus, $\tau_1 \geq r/\theta$. Now, let $\Delta \neq \emptyset$ and $P(\Delta) > 0$. Let also $\rho = \lfloor r/\theta \rfloor$.

$$\mathbb{E}[W \mid D = \Delta] = \mathbb{E}[W \mid D \neq \emptyset] = \mathbb{E}[\tau_2 + 1 \mid \tau_1 \geq \tau_2]$$

Since $\mathbb{E}[\tau_2 \mid \tau_2 \leq \min\{\rho, \tau_1\}] \leq \rho < \mathbb{E}[\tau_2 \mid \rho < \tau_2 \leq \tau_1]$, and $\tau_2 \leq \rho$ implies $\tau_2 \leq \tau_1$,

$$\mathbb{E}[W \mid D = \Delta] \geq 1 + \mathbb{E}[\tau_2 \mid \tau_2 \leq \rho]$$

Also,

$$\begin{aligned}\mathbb{E}[\tau_2 \mid \tau_2 \leq \rho] &= \sum_{j=1}^{\rho} j \Pr[\tau_2 = j \mid \tau_2 \leq \rho] \\ &= \sum_{j=1}^{\rho} j \frac{\Pr[\tau_2 = j]}{\Pr[\tau_2 \leq \rho]} = \sum_{j=1}^{\rho} j \frac{\pi q^{j-1}}{1 - q^{\rho}}\end{aligned}$$

where $q = 1 - \pi$. After some computations we obtain

$$\mathbb{E}[\tau_2 \mid \tau_2 \leq \rho] = \frac{1 - q^{\rho}(1 + \pi\rho)}{\pi(1 - q^{\rho})}$$

Since $(1 - q^{\rho}) \leq 1$ and $q^{\rho} \leq e^{-\pi\rho}$, $\mathbb{E}[\tau_2 \mid \tau_2 \leq \rho] \geq (1 - e^{-\pi\rho}(1 + \pi\rho))/\pi \geq 0.15/\pi$, for $\pi\rho \geq 0.7$. Also, it is $\mathbb{E}[\tau_2 \mid \tau_2 \leq \rho] = (1 + \pi\rho - \pi\rho/(1 - q^{\rho}))/\pi$, and since $q^{\rho} \leq 1 - \pi\rho + (\pi\rho)^2/2$, $\mathbb{E}[\tau_2 \mid \tau_2 \leq \rho] \geq \rho(1 - \pi\rho)/(2 - \pi\rho) \geq 0.2\rho$, for $\pi\rho \leq 0.7$. Therefore, $\mathbb{E}[W \mid D = \Delta] \geq 1 + 0.15 \cdot \min\{\rho, 1/\pi\}$. Part (b) of the fact follows from this result and the bound for $\mathbb{E}[W \mid D = \emptyset]$ we showed earlier. □

We now describe how G' and $W_0, W_1, \ldots, W_{m'-1}$ are constructed in terms of G. We begin with an informal exposition. The construction is similar, in principle, to that in the proof of Lemma 3. Let $\hat{R}$ be the subgraph of the routing tree of G induced by the nodes in $\langle 0, m\rangle$. ($\hat{R}$ is itself a tree.) Let also R' denote the routing tree of G'. We construct the pairs (Δ'_u, W_u) inductively by considering each node $u = 0, 1, \ldots, m'-1$ in turn. For each u we distinguish two cases depending on whether u is a leaf of R' or not. If u is a leaf, we draw (Δ'_u, W_u) from $\hat{P}'$ independently of G and other random choices we make. With each u that is not a leaf of R' we associate a node C_u of $\hat{R}$. As in the construction in the proof of Lemma 3, $C_0 = 0$ and if $u > 0$, C_u is determined at the same time as the delta-set of the parent of u in R'. Roughly, if u is not a leaf of R', Δ'_u and W_u are determined from Δ_{Y_u} and $U(G, C_u, Y_u)$, respectively, where Y_u is a descendant of C_u in $\hat{R}$. (We give the details below.) Also, with each (non-leaf) child $u+\delta$ of u in R' we associate a distinct child $C_{u+\delta}$ of Y_u in $\hat{R}$, such that $C_{u+\delta} \approx Y_u + \delta r'$.

More precisely, let A_k, for $k < m$, be the r-ancestor of node k in G. Let also $\Xi_k = k$ if k is an r-significant node of G (i.e., if $A_k = k$), and $\Xi_k = A_k + r'$ otherwise. Finally, let H_u, for $u < m'$, be the span of u in the routing tree of $(\Delta'_0, \Delta'_1, \ldots, \Delta'_{u-1})$. For each $u \in [0..m'-1]$,

- if $H_u = 1$, the value of pair $(\Delta'_u,\ W_u)$ is chosen independently at random from $\hat{P}'$.
- if $H_u > 1$,

$$\Delta'_u = \Phi(\Delta_{Y_u}), \quad W_u = U(G, C_u, Y_u) + 1$$

  where:

  - Y_u is the first node in the rightmost path from C_u in $\hat{R}$, such that: (i) a child of Y_u in $\hat{R}$ is $\geq C_u + r$ **or** (ii) $\Delta_{Y_u} \cap [\theta..m-1] \neq \emptyset$
  - $C_0 = 0$, and if $u > 0$ and v is u's parent in the routing tree of $(\Delta'_0, \Delta'_1, \ldots, \Delta'_{u-1})$,

$$C_u = Y_v + \max\big\{d \in \{1\} \cup \Delta_{Y_v} \cap \langle 0, m\rangle : (Y_v + d) - \Xi_{C_v} \leq (u - v) \cdot r'\big\}$$

The next fact gives more insight into the construction above; it is the analogous of Facts 3.1 and 3.2. (The proof is omitted.) Let $\hat{S}_k$, for $k < m$, be the size of the subtree of $\hat{R}$ rooted at k. Recall also that H_u, for $u < m'$, is equal to the size of the subtree of R' rooted at u.

FACT 4.1.2. *For each $u \in [0..m'-1]$, if $H_u > 1$ then*

(a) $0 \le C_u \le \Xi_{C_u} \le ur'$ and $C_u \le Y_u < C_u + r$

(b) $\Xi_{C_u} + H_u r' \le C_u + \hat{S}_{C_u} = Y_u + \hat{S}_{Y_u}$

(c) if $u < u' < u + H_u$ and $H_{u'} > 1$ then $Y_u < C_{u'} < Y_u + \hat{S}_{Y_u}$;
if $u + H_u \le u' < m'$ and $H_{u'} > 1$ then $C_{u'} \ge Y_u + \hat{S}_{Y_u}$

We now show that G' and $W_0, W_1, \ldots, W_{m'-1}$ have the desired distribution. As in the proof of Lemma 3, we think of the construction as a random m'-step process, where in step u we determine the value of (Δ'_u, W_u). We assume that the value of each Δ_k is generated right before it is about to be used for the first time — not earlier. Let $\mathcal{D}_u$ be the subset of G's nodes consisting of all k for which Δ_k is generated in steps 0 up to $u-1$. Observe that $\mathcal{D}_u$ consists of the nodes in the routing paths of G from C_v to Y_v, for all $v < u$. So, Fact 4.1.2(c) implies that for all $k \ge C_u$, $k \notin \mathcal{D}_u$; thus, the delta-set of each node in the path from C_u to Y_u is generated independently in step u. It is now straightforward to verify that (Δ'_u, W_u) is distributed according to $\hat{P}'$, and it is independent of choices made in previous steps.

Next we bound the expected "weighted" routing cost in G' to the average destination in terms of the corresponding "unweighted" quantity. For $u, v < m'$, let $U_{\mathrm{w}}(u,v) = \sum_{u' \in \vartheta_{u,v}} W_{u'}$, where $\vartheta_{u,v}$ is the set of all the nodes in the routing path from u to v in G', *excluding* the last node v. Let also $\bar{U}_{\mathrm{w}} = (1/m') \sum_{u=0}^{m'-1} U_{\mathrm{w}}(0,u)$. (So, $U_{\mathrm{w}}(u,v)$ and $\bar{U}_{\mathrm{w}}$ are the weighted versions of $U(G',u,v)$ and $\bar{U}(G')$, respectively.) We have

$$\mathbb{E}[U_{\mathrm{w}}(0,u) \mid G'] = \sum_{v \in \vartheta_{0,u}} \mathbb{E}[W_v \mid G'] = \sum_{v \in \vartheta_{0,u}} \mathbb{E}[W_v \mid \Delta'_v]$$
$$= U(G',0,u) \cdot \mathbb{E}[W \mid D]$$

The second equality holds because of the independence of the (Δ'_i, W_i) pairs; the last holds because each of these pairs is distributed identically to (D, W), and $|\vartheta_{0,u}| = U(G',0,u)$. From the above result and Fact 4.1.1(b) we obtain that $\mathbb{E}[U_{\mathrm{w}}(0,u) \mid G'] \ge c_1 \min\{r/\theta, 1/\pi\} \cdot U(G',0,u)$. Taking the average over all $u < m'$, and the expectation over G', we get

$$\mathbb{E}[\bar{U}_{\mathrm{w}}] \ge c_1 \min\{r/\theta, 1/\pi\} \cdot \mathbb{E}[\bar{U}(G')]$$
$$\ge c_1 \min\{r/\theta, 1/\pi\} \cdot T(m', \ell')$$

By Lemma 1 and Fact 4.1.1(a), $T(m',\ell') \ge T(m', 4\lambda/\pi)$. By Lemma 3, we can omit factor 4 in the last inequality at the cost of a constant factor, and by Lemma 2, we can replace $m' = \lfloor m/(r+\theta) \rfloor$ with $\lfloor m/r \rfloor \le 2m'$ again at the cost of a constant factor; thus, $T(m',\ell') \ge c_2 T(\lfloor m/r \rfloor, \lambda/\pi)$. Finally, by Lemma 3 and the fact that $T(i) \ge 2/3$ for $i \ge 3$, $T(\lfloor m/r \rfloor, \lambda/\pi) \ge \max\{2/3,\ (\pi/\lambda) T(\lfloor m/r \rfloor\}$. Therefore,

$$\mathbb{E}[\bar{U}_{\mathrm{w}}] \ge c_3 \min\{r/\theta,\ (1/\lambda) T(\lfloor m/r \rfloor)\} \qquad (7)$$

To complete the proof, it remains to bound $\mathbb{E}[U_Z]$ in terms of $\mathbb{E}[\bar{U}_{\mathrm{w}}]$. We show that

$$\sum_{k < r'm'} U(G,0,A_k) \ge r' \sum_{u < m'} U_{\mathrm{w}}(0,u) \qquad (8)$$

We do so by proving that the following predicate holds for all $h \in [2..m']$ (for fixed G): "for all $u \in [0..m'-1]$ such that $H_u = h$,

$$\sum_{v \in \langle u, H_u \rangle} U_{\mathrm{w}}(u,v) \le (1/r') \sum_{k \in \langle \Xi_{C_u}, r'H_u \rangle} U(G, C_u, A_k)\text{"}$$

(The proof is by induction on h; the details of this proof are omitted.) (8) follows then by taking $h = m'$. From (8) (by dividing both sides by m, and taking the expectation over G) we obtain that $\mathbb{E}[U_Z] \ge c_4\, \mathbb{E}[\bar{U}_{\mathrm{w}}]$. Substituting $\mathbb{E}[\bar{U}_{\mathrm{w}}]$ with the right side of (7) we obtain the desired bound for $\mathbb{E}[U_Z]$.

3.5.2 Proof of Fact 4.1(b)

Let S be the set of the r-significant nodes of G that belong to $\langle 0, m \rangle$. Let also $\hat{S} = \{u \in S : D_u = r\}$ and $\check{S} = \{u \in S : D_u < \eta\}$, where D_u denotes the number of the r-descendants of u in G that are in $\langle 0, m \rangle$. We have,

$$\mathbb{Pr}[M < \eta] = \mathbb{E}\left[\mathbb{Pr}[M < \eta \mid G]\right] = \mathbb{E}\left[\frac{1}{m} \sum_{u \in \check{S}} D_u\right]$$
$$\le \mathbb{E}\left[\frac{1}{m} |\check{S}| \cdot (\eta - 1)\right] \le \frac{\eta - 1}{m}\, \mathbb{E}[|S| - |\hat{S}|]$$

Note that except for node 0, every node in S is an r-successor of some node in $\hat{S}$. Also if $u \in \hat{S}$, the number of u's r-successors is less or equal to the total number N_u of long-range contacts of the nodes in the r-border of u, plus 1 (for the successor of $u + r - 1$). From the last two observations, $|S| \le 1 + \sum_{u \in \hat{S}} (N_u + 1) = 1 + |\hat{S}| + \sum_{u \in \hat{S}} N_u$. So,

$$\mathbb{Pr}[M < \eta] \le \frac{\eta - 1}{m}\Big(1 + \mathbb{E}\Big[\sum_{u \in \hat{S}} N_u\Big]\Big) \qquad (9)$$

Next, we bound $\mathbb{E}[\sum_{u \in \hat{S}} N_u]$. Let R^t, for $0 \le t \le m$, be the routing tree of $(\Delta_0, \Delta_1, \ldots, \Delta_{t-1})$. (Note that R^t completely determines which of the nodes in $[0..t]$ are in $\hat{S}$.)

FACT 4.1.3. $\mathbb{E}[N_u \mid R^u, u \in \hat{S}] = \mathbb{E}[U_{r-1}] + 1$.

PROOF. Assuming $u \in \hat{S}$, let τ be the length of the r-border of u. Also for each $j \ge 0$, let L_j be the number of long-range contacts of the $(j+1)^{\text{th}}$ node along this path, if $j \le \tau$; if $j > \tau$ the value of L_j is drawn independently at random from the distribution of $|\Delta_k|$. Finally, let A be a shorthand for "$R^u, u \in \hat{S}$."

$$\begin{aligned}
&\mathbb{E}[N_u \mid A] \\
&= \sum_i \mathbb{E}[N_u \mid A, \tau = i] \cdot \mathbb{Pr}[\tau = i \mid A] \\
&= \sum_i \sum_{j=0}^{i} \mathbb{E}[L_j \mid A, \tau = i] \cdot \mathbb{Pr}[\tau = i \mid A] \\
&= \sum_{j \ge 0} \Big(\mathbb{E}[L_j \mid A] - \mathbb{E}[L_j \mid A, \tau < j] \cdot \mathbb{Pr}[\tau < j \mid A] \Big) \\
&= \sum_{j \ge 0} (1 - \mathbb{Pr}[\tau < j \mid A]) = 1 + \sum_{j \ge 0} (1 - \mathbb{Pr}[\tau \le j \mid A]) \\
&= 1 + \mathbb{E}[\tau \mid A] = 1 + \mathbb{E}[U_{r-1}]
\end{aligned}$$

The forth equality holds because, conditioned on either A or $A \cap \{\tau < j\}$, L_j is distributed like $|\Delta_k|$, thus, $\mathbb{E}[L_j \mid A] = \mathbb{E}[L_j \mid A, \tau < j] = 1$. □

Let Z_i, for $1 \le i \le m+1$, be the i^{th} largest node in $\hat{S}$ if $i \le |\hat{S}|$, or $Z_i = m$ otherwise. Let also $N_m = 0$. For each $t \in [0..m]$, we define

$$X_t = \sum_{j \in [1..t]} N_{Z_j} + \frac{m - Z_{t+1}}{r} \cdot (1 + \mathbb{E}[U_{r-1}])$$

Using Fact 4.1.3, it is straightforward to verify that $\mathbb{E}[X_{t+1} - X_t \mid R^u, Z_{t+1} = u] \le 0$. Therefore, $\mathbb{E}[X_{t+1} - X_t] \le 0$, which implies $\mathbb{E}[X_m] \le X_0$, or, equivalently, $\mathbb{E}[\sum_{u \in \hat{S}} N_u] \le (m/r)(1 + \mathbb{E}[U_{r-1}])$. This, together with (9), yield the desired bound for $\Pr[M \ge \eta]$.

4. CONCLUDING REMARKS AND FUTURE WORK

In this paper, we have shown that the expected number of steps for greedy routing in uniform ring-based random graphs of n nodes with expected ℓ long-range contacts per node is $\Omega((\log^2 n)/\ell a^{\log^* n})$. This improves a lower bound by Aspnes *et al.*, and proves that greedy routing in this class of graphs is suboptimal even when $\ell = \Theta(\log n)$, a case of practical interest.

Our lower bound is very close to the upper bound of $O((\log^2 n)/\ell)$ that greedy routing achieves in Kleinberg's small-world networks, a particular instance of uniform ring-based random graphs. Our analysis suggests that the optimal distribution P for choosing long-range contacts has structural properties similar to those of the distribution used in Kleinberg's construction. We conjecture that our lower bound can be improved to $\Omega((\log^2 n)/\ell)$, i.e., that Kleinberg's construction is in fact (asymptotically) optimal for greedy routing in uniform ring-based random graphs.

In this paper we have focused on unidirectional greedy routing, where the distance from node u to node v is the number of edges along the ring in, say, clockwise direction from u to v. In bidirectional greedy routing, the distance between two nodes is the minimum number of edges between them in either clockwise or counterclockwise direction. In most actual designs, both versions of greedy routing give (asymptotically) the same results. We conjecture that the same asymptotic bounds apply to both versions.

The model of uniform ring-based random graphs naturally generalizes to more than one dimensions, by using the d-dimensional torus (or grid) as a base graph instead of the ring. Kleinberg's construction also generalizes to higher dimensions resulting in the same $O((\log^2 n)/\ell)$ upper bound for greedy routing. It is interesting to investigate whether the use of additional dimensions improves the performance of greedy routing or if a lower bound similar to that for the one-dimensional case applies.

5. ACKNOWLEDGMENTS

We thank the anonymous referees for their comments.

6. REFERENCES

[1] I. Abraham, D. Malkhi, and G. S. Manku. Papillon: greedy routing in rings, `http://arxiv.org/abs/cs/0507034`, 2005. (See also *Proc. 19th Intl. Symp. on Distributed Computing (DISC)*, pages 514–515, 2005.)

[2] J. Aspnes, Z. Diamadi, and G. Shah. Fault-tolerant routing in peer-to-peer systems. In *Proc. 21st ACM Symp. on Principles of Distributed Computing (PODC)*, pages 223–232, 2002.

[3] L. Barrière, P. Fraigniaud, E. Kranakis, and D. Krizanc. Efficient routing in networks with long range contacts. In *Proc. 15th Intl. Symp. on Distributed Computing (DISC)*, pages 270–284, 2001.

[4] M. Flammini, L. Moscardelli, A. Navarra, and S. Pérennes. Asymptotically optimal solutions for small world graphs. In *Proc. 19th Intl. Symp. on Distributed Computing (DISC)*, pages 414–428, 2005.

[5] P. Ganesan and G. S. Manku. Optimal routing in Chord. In *Proc. 15th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 176–185, 2004.

[6] K. Gummadi, R. Gummadi, S. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica. The impact of DHT routing geometry on resilience and proximity. In *Proc. ACM SIGCOMM 2003*, pages 381–394, 2003.

[7] M. F. Kaashoek and D. R. Karger. Koorde: A simple degree-optimal hash table. In *Proc. 2nd Intl. Workshop on Peer-to-Peer Systems (IPTPS)*, pages 98–107, 2003.

[8] J. Kleinberg. The small-world phenomenon: an algorithm perspective. In *Proc. 32nd ACM Symp. on Theory of Computing (STOC)*, pages 163–170, 2000.

[9] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proc. Intl. Congress of Mathematicians (ICM)*, 2006.

[10] D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: a scalable and dynamic emulation of the butterfly. In *Proc. 21st ACM Symp. on Principles of Distributed Computing (PODC)*, pages 183–192, 2002.

[11] G. S. Manku. Routing networks for DHTs. In *Proc. 22nd ACM Symp. on Principles of Distributed Computing (PODC)*, pages 133–142, 2003.

[12] G. S. Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in a small world. In *Proc. 4th USENIX Symp. on Internet Technologies and Systems (USITS)*, pages 127–140, 2003.

[13] G. S. Manku, M. Naor, and U. Wieder. Know thy neighbor's neighbor: the power of lookahead in randomized P2P networks. In *Proc. 36th ACM Symp. on Theory of Computing (STOC)*, pages 54–63, 2004.

[14] M. Naor and U. Wieder. Novel architectures for P2P applications: the continuous-discrete approach. In *Proc. 15th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 50–59, 2003.

[15] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for Internet applications. *IEEE/ACM Transactions on Networking*, 11(1):17–32, 2003.

[16] J. Xu. On the fundamental tradeoffs between routing table size and network diameter in peer-to-peer networks. In *Proc. IEEE INFOCOM 2003*, 2003.

[17] H. Zhang, A. Goel, and R. Govindan. Incrementally improving lookup latency in distributed hash table systems. In *Proc. ACM SIGMETRICS 2003*, pages 114–125, 2003.