A Contribution to Analyzing and Enhancing Darknet Routing

Stefanie Roos Thorsten Strufe
TU Darmstadt
<lastname>@cs.tu-darmstadt.de

***Abstract*—Routing in Darknets, membership concealing overlays for pseudonymous communication, like for instance Freenet, is insufficiently analyzed, barely understood, and highly inefficient. These systems at higher performance are promising privacy preserving solutions for social applications. This paper contributes a realistic analytical model and a novel routing algorithm with provable polylog expected routing length. Using the model, we additionally prove that this can not be achieved by Freenet's routing. Simulations support that our proposed algorithm achieves a better performance than Freenet for realistic network sizes.**

I. Introduction

Darknets represent a highly promising system class to provide a communication substrate for decentralized, social applications. Such overlays offer secure and private communication, desired by various social applications, though the term Darknet is commonly associated with sharing of illegal content. They implement messaging and content publication, which can be used to model all usual social communication functions, at very high confidentiality. The principle of only having connections to trusted contacts is an intuitive solution for systems dealing with sensitive and private information, such as social network profiles, and censorship-resilient publication of opinions and criticism. Achieving both sender and receiver anonymity, as well as membership-concealment, they offer high privacy guarantees, yet, their current primary drawback is their lack of performance. This deficiency probably is the foremost reason for their limited employment as censorship-resistant publication systems, as well.

The foundations of Darknets to achieve privacy are the main cause for their low performance, too. They rely on 1) permitting connections between nodes only if the respective individuals running them share a mutual trust relationship to hide the participation from any untrusted, potentially malicious party; 2) applying source rewriting on all forwarded messages to conceal their initiator and path; and 3) encrypting requests and content to achieve confidentiality. Given (1), the Darknet topologies resemble the scale-free graphs of the underlying social networks.

Considering prior results [1], it is assumed that a greedy routing algorithm can be found that converges in an expected polylog number of hops on those networks. Due to the connectivity restrictions, it is a difficult problem to implement such an efficient routing, though.

Existing proposals, such as Turtle [2] and OneSwarm [3], use flooding, which scales linearly in the network size at best. Creating a DHT-like system on a Darknet topology, by establishing multi-hop tunnels to construct the necessary neighborhood for a greedy routing, has been proposed in [4], [5], as a more efficient solution. However, constructing and maintaining the tunnels causes high state and maintenance costs. Freenet, the only actually deployed Darknet using a deterministic routing algorithm, adjusts the node identifiers to the fixed topology [6]. These node identifiers are then used to enable deterministic routing. For routing, Clarke et el. propose a *distance-directed depth first search* (D^2*-DFS*), to address the fact that a perfect embedding can not be achieved.

In prior work, we proposed a simplified model and a class of Darknet routing algorithms [7]. In this paper, we contribute a formal model to analyze Darknets, which extends Kleinbergs small-world model, but more accurately reflects the imprecision of the embedding, the bidirectional links of mutual trust relationships, and finally the scale-free character of social graphs. Using this model, we are able to prove that the current routing in Freenet indeed does not achieve an expected polylog routing length, whereas our routing algorithm *NextBestOnce* is shown to meet this requirement. Extensive simulation studies, however, are inconclusive: though supporting the polylog performance of *NextBestOnce*, they show that Freenet's routing algorithm actually achieves quite low absolute path lengths, and outperforms *NextBestOnce* for small network sizes.

In the remainder of this paper, we first explain the foundations of Freenet in higher detail and formalize model and problem description. We subsequently analyze the performance of Freenet's routing algorithm, introduce our routing algorithm *NextBestOnce*, and prove its performance to be polylog. The results of extensive simulations are presented thereafter, and we close our paper with a conclusion.

II. Models and Problem Definition

In this section, we first briefly discuss some related work on Darknet modelling, before presenting our model.

A. Background

A Darknet is an overlay network, in which connections correspond to a mutual trust relationship between the respective participants. By this, Darknet topologies are social graphs, induced by real-world relationships of individuals.

Social graphs are commonly assumed to be:

- scale-free, i.e. the probability that the degree D of a node is d is given by $P(D = d) \propto \frac{1}{d^\alpha}$ for some $\alpha \in [2, 3]$

- small-world, i.e. the diameter of the graph is logarithmic to the network size.

A Darknet topology model hence has to include these characteristics of social networks. It additionally has to include a namespace, i.e. a mapping from nodes to identifiers in a metric space, to permit modeling the routing. Freenet chooses the small-world topology model by Jon Kleinberg [1] as an analytic foundation: Nodes are arranged in a multidimensional grid, edges exist between nodes that are closest to each other, and each node has one directed edge to a neighbor chosen with a probability anti-proportional to the distance.

In difference to such a generative model, or to conventional peer-to-peer systems for that matter, nodes in Darknets can not establish connections to the nodes that are closest in the namespace, due to the restriction to connect to trusted nodes only. (Throughout the paper, the distance of nodes refers to the distance in the namespace, rather than the hop distance.) For that reason, it is not easily possible to create the common lattice structure (with additional links for performance gains) to facilitate straight forward greedy routing along the namespace. A routing structure can only be *approximated* by assigning suitable identifiers to the nodes, thus finding a mapping of the nodes into a metric space, which commonly is termed *embedding*. Such an approximation, however, is not well reflected by Kleinberg's small-world model. Rather than the lattice structure, our model assumes that nodes are connected to some nodes in their vicinity, but not necessarily to the closest nodes.

B. Darknet Model $\mathcal{D}(n, d, C, L)$

We extend Kleinberg's model by the two parameters C, the maximal distance to the closest neighbor over all nodes, and the random variable L defining the assumed degree distribution. Each node $v = (v_1, ..., v_d)$ has hence short-range links to neighbors (with higher and lower ID) in each direction: $a^v_1, ..., a^v_d, b^v_1, ..., b^v_d$. Here a^v_j is chosen from the set

$$\begin{aligned} A^v_j = \{u = (u_1, ..., u_d) \in V : u_i = v_i \text{ for } i \neq j, \\ 1 \leq \min\{u_j - v_j, n + u_j - v_j\} \leq C\}. \end{aligned} \quad (1)$$

Analogously, b^v_j is chosen from

$$\begin{aligned} B^v_j = \{u = (u_1, ..., u_d) \in V : u_i = v_i \text{ for } i \neq j, \\ 1 \leq \min\{v_j - u_j, n + v_j - u_j\} \leq C\}. \end{aligned} \quad (2)$$

In favor of a more coherent presentation, we assume w.l.o.g. that the neighbors are chosen uniformly at random from A^v_i and B^v_i. The same results can be derived if each element in the set is chosen with arbitrary, non-zero probability.

In addition to the short-range links, long-range links are chosen in a two step process:

1) choose a label $l_v \in \mathbb{N}$, distributed according to L, for each node $v \in V$
2) connect nodes $u, v \in V$ with probability

$$P(l(u,v)|l_u = d_1, l_v = d_2) = 1 - e^{-\frac{d_1 d_2}{dist(u,v)^d \gamma}} \quad (3)$$

Basic calculations show that for a node $v \in V$, the expected degree given the label l_v is $\mathbb{E}(D_v|l_v) = \Theta(l_v)$. So, a scale-free distribution L leads to a scale-free degree distribution. Additionally, the probability that two arbitrary nodes $u, v \in V$ are adjacent is

$$P(l(u,v)) = \Theta\left(\frac{1}{dist(u,v)^d \cdot \gamma}\right), \quad (4)$$

corresponding to the original (directed) small-world model by Kleinberg.

In the following, two basic results are given. They are essential for deriving both the lower and upper bounds on the routing length in Sections III.

Lemma 2.1: Denote by $l(u,v)$ the fact that u and v are linked via a long-range link. Two arbitrary nodes u, v are long-range neighbors with probability

$$P(l(u,v)) = \Theta\left(\frac{1}{dist(u,v) \cdot \log n}\right). \quad (5)$$

The probability that the distance between u and v exceeds $\sqrt{n}$ is given by

$$p_l = P(dist(u,v) > \sqrt{n}|l(u,v)) = \Theta(1). \quad (6)$$

Lemma 2.2: With probability at most $\frac{|P|}{\sqrt{n}}$ a node v with distance at least $\sqrt{n}$ to t is contained in the routing path P.

The proofs for both lemmata are omitted for lack of space, but provided in [8].

III. Performance of D^2-*DFS*

In this section we analyze the performance of D^2-*DFS* in the context of $\mathcal{D}(n, 1, C, L)$, restricted to $d = 1$ since Freenet uses a single dimension. D^2-*DFS* works as follows: Each node chooses the neighbor closest to the destination that is not known to have received the message before (i.e. neither predecessor nor previously contacted neighbors) as a next hop, if such a node exists. In case that a receiving node has previously received the message or no further neighbors are available to contact, the message is backtracked to the predecessor. The only requirement with regard to the degree distribution is that the degree of a node is bounded by a constant T with probability r, meaning that the degree of a certain percentage of nodes does not increase with the network size.

The performance is given by the expected routing length. For two distinct nodes s, t the routing length is denoted $R^{DFS}(s,t)$, and the expected routing length for the whole graph is given as $\mathbb{E}(R^{DFS}) = \frac{1}{n(n-1)} \sum_{s \neq t} \mathbb{E}(R^{DFS}(s,t))$.

Theorem 3.1: Let L be such that the degree D_u of node u is bounded by a constant $T \in \mathbb{N}$ with constant probability $r \in \mathbb{R}_+$, i.e. $P(D_u \leq T) \geq r > 0$, and $C > 2$. Then D^2-*DFS* does not have polylogarithmic expected routing length, i.e. for any $\rho > 0$:

$$\frac{1}{n(n-1)} \sum_{s \neq t \in V} \mathbb{E}(R^{DFS}(s,t)) = \Omega\left(\log^\rho n\right) \quad (7)$$

The proof is split into three lemmata. The first one, Lemma 3.2, shows that a long-range link is used with constant probability, even though the message is already very close to the target t. The message can afterwards only reach t during backtracking or via a different long-range link. Furthermore, it is shown that the average number of nodes for which this might happen grows linearly with the network size. In the remaining section, it is proven that the probability to find such a long-range link within $M \log^\rho n$ is negligible (Lemma 3.3). The same holds for returning to the node by backtracking (Lemma 3.4). Hence, the routing length $R^{DFS}(s,t)$ exceeds $M \log^\rho n$ with probability $p > 0$ for any s,t with $dist(s,t) > S_t$ where S_t denotes the local neighborhood of t. So Theorem 3.1 follows because $E(R^{DFS}(s,t)) > \frac{M}{p} \log^\rho n$, with M chosen arbitrarily.

For any node $u \in V$, let u_m be the node with ID $id(t) + m \ mod \ n$. Given the target t, we consider a set $S_t = \{t_{-m_1}, \cdots, t, \cdots, t_{m_2}\}$ for some constants m_1, m_2 and show that with constant probability a message is forwarded in such a way that t can only be reached via a long-range link.

Lemma 3.2: For a set $S_t = \{t_{-m_1}, ..., t, ..., t_{m_2}\}$, containing the neighborhood of t, the probability q_C that all nodes in $S'_t := S_t \setminus \{t_{-m_1}, t_{m_2}\}$ have only short-range links to nodes in S_t depends only on C. In such a case, D^2-*DFS* marks both t_{-m_1} as well as t_{m_2} with constant probability q_O before forwarding the message away from the target to a node in $V \setminus S_t$ connected through a long-range link. Consequently, with constant probability, t can only be reached via a long-range link or during the backtracking phase.

Proof: A lower bound on the probability for an adverse short-range link selection can be given for any S_t. By example, we show that with constant probability q_O, a message is forwarded via a long-range link with the described result.

The probability that the short-range neighbor $v_u^\pm$ of a node u, i.e. the neighbor with the higher (v_u^+) respectively lower identifier (v_u^-), is contained in any subset $H \subset V$ is given as the ratio between the nodes in H that can be chosen as $v_u^\pm$ and all C nodes that can be chosen as $v_u^\pm$, i.e.

$$P(v_u^\pm \in H) = \frac{|H \cap \{u_{\pm 1}, \cdots, u_{\pm C}\}|}{C}. \tag{8}$$

The probability q_C of having no short-range links between S'_t and $V \setminus S_t$ is computed as:

$$q_C \geq \left(\prod_{i=1}^{C-1} \frac{i}{C} \right)^4 \tag{9}$$

The inequality (equality holds if $|S_t| \geq 2C$) follows since at most $4(C-1)$ nodes can have a neighbor in the other set, with a probability depending on their distance to the set. For example, the node t_{-m1-1} chooses a neighbor within $\{t_{-m1}, t_{-m1+1}, ..., t_{-m1+C-1}\}$, so with probability $\frac{1}{C}$ the neighbor is not in S'_t, namely if it is t_{-m1}. The same holds for t_{m2+1} and similarly for t_{m1+1}, t_{m2-1} when replacing S'_t with $V \setminus S_t$. For t_{m2-2} the chance to choose a node not in S'_t is then $\frac{2}{C}$, and so on.

By this, we have shown that the lower bound on q_C depends only on C, not on n and S_t.

As a result of such a short-range link, the only possibility for a node $s \in V \setminus S_t$ to route to a node $t \in S$ is to take a path containing t_{-m_1}, t_{m_2} or using a long-range link to a node in S_t. If both t_{-m_1}, t_{m_2} are marked, backtracking has to be used in case no long-range link is found.

Figure 1 illustrates an example of such a path. The case $S_t = \{t_{-4}...t_{+3}\}$ is considered. Starting from t_{+3} gives the partial path $t_{+3}, t_{+1}, t_{-1}, t_{-4}, t_{-2}$. t_{-2}'s only short-range neighbor is t_{+1}, which is already on the routing path. Hence, the message is forwarded using a long-range link. The probability that t_{-2} has at least one long-range link exceeds 0 for any non-trivial degree distribution. It is easy to show that with probability $q_O = \Theta(1)$, t_{+3} is the first node in S_t that is contacted. ■

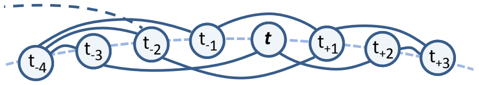

Fig. 1: Exemplary adverse connectivity for D^2-*DFS*

The destination t can only be reached by long-range links and by backtracking if the message is forwarded as in Lemma 3.2.

Next, we bound the probability for this to happen within $M \log^\rho n$ steps, starting with the case of finding a long-range link to a node in S_t.

Lemma 3.3: The probability that during D^2-*DFS* no node in S_t is chosen via a long-range link is at least $q_R = 1 - \frac{1}{2} p_l^{|S_t|T} r^{|S_t|} > 0$ for n big enough and p_l as in Eq. 6, assuming that the path length is maximally $M \log^\rho n$.

Proof: Considering the complementary event of finding a long-range link to a node in S_t, with probability $r^{|S_t|}$ there are at most $|S_t|T$ long-range links into S_t. All of these links lead to nodes in at least distance $\sqrt{n}$ with probability at least $p_l^{|S_t|T}$, as defined in Eq. 6. The probability that $|S_t|T$ nodes with distance at least $\sqrt{n}$ to t are not contained in a path of length $M \log^\rho n$ is at least

$$\left(1 - \frac{M \log^\rho n}{\sqrt{n}}\right)^{|S_t|T} > \frac{1}{2}$$

for n big enough by Lemma 2.2. So

$$q_R = 1 - \frac{1}{2} p_l^{|S_t|T} r^{|S_t|}$$

■

Lemma 3.4: With probability of at least

$$q_B = \frac{1}{2} r \cdot p_l^T > 0,$$

D^2-*DFS* does not backtrack to a local node v_c, the node that forwarded the message to a node connected through a long-range link $v_f \notin S_t$, within $M \log^\rho n$ steps.

Proof: The probability that all long-range links of v_c have length at least $\sqrt{n}$ is bounded from below by $r \cdot p_l^T$ using Lemma 2.1. D^2-*DFS* only considers v_c a second time, when all nodes reachable from v_f have been visited, without contacting any node already on the routing path. Consider the $M \log^\rho n$

nodes reachable from v following short-range links opposite to t. By Lemma 2.2, each of these nodes is on the path with a probability of at most $\frac{M\log^{\rho} n}{\sqrt{n}}$. Hence, the probability that none of them is on the path is

$$\left(1-\frac{M\log^{\rho} n}{\sqrt{n}}\right)^{M\log^{\rho} n} > \frac{1}{2}$$

for n big enough. So with probability of at least $q_B = \frac{1}{2}r \cdot p_l^T$, $M\log^{\rho} n$ nodes need to be considered before considering v_c a second time. ■

The proof of Theorem 3.1 merely combines Lemma 3.2, 3.3 and 3.4.

Proof: Note that showing that D^2-*DFS* needs more than $M\log^{\rho} n$ steps is equivalent to showing that it needs $\frac{M}{p}\log^{\rho} n$ steps for $p > 0$, since M can be any constant. The probability that the short-range links are chosen in an adverse way is q_C. With constant probability q_O, the nodes are then visited in an order, so that the message is forwarded to a long-range neighbor of v_c. The probability that at least $M\log^{\rho} n$ are needed before backtracking is q_B by Lemma 3.4. Similarly, the probability of not contacting a node in S_t using a long-range link is q_R. Combining this, the probability that D^2-*DFS* needs at least $M\log^{\rho} n$ is at least $q_C \cdot q_R \cdot q_B \cdot q_{+3}$. So,

$$\mathbb{E}(R^{DFS}(s,t)) > q_C \cdot q_O \cdot q_R \cdot q_B \cdot M\log^{\rho} n$$

Hence, for any $M, \rho > 0$, we have $E(R^{DFS}(s,t)) > M\log^{\rho} n$. Because s and t are arbitrary nodes with $dist(s,t) > |S_t|$, the average expected routing length over all nodes is bounded from below:

$$\begin{aligned}
&\frac{1}{n(n-1)}\sum_{s\neq t\in V}\mathbb{E}(R^{DFS}(s,t)) \\
\geq &\frac{1}{n(n-1)}\sum_{s\neq t\in V, dist(s,t)>|S_t|}\mathbb{E}(R^{DFS}(s,t)) \qquad (10)\\
\geq &\frac{n(n-|S_t|)}{n(n-1)}M\log^{\rho} n = \Omega(\log^{\rho} n)
\end{aligned}$$

■

Theorem 3.1 does not give an exact bound for D^2-*DFS*. It proves that even though short paths exist, D^2-*DFS* does not achieve polylog routing length, if the applied embedding does not achieve that each local link has a maximum distance of 2.

IV. NextBestOnce

D^2-*DFS* has two drawbacks that can increase the routing length. The first one is that nodes have to be contacted to check if the message has already passed them. This results in a message overhead, that is not necessary in case nodes are aware if their neighbors have already seen the message. The second drawback is that a node on the path always contacts the neighbor that has not yet seen the message and is closest to the destination. As we have seen in Section III, this might cause the message to be passed along a long-range link away from the destination, and degrade the routing performance.

The first issue can easily be solved by including information about *marked* nodes, i.e. nodes that should not be contacted again. The second drawback is harder to resolve. The main idea of *NextBestOnce* is to forward the message to the neighbor closest to the destination, possibly passing nodes several times. Nevertheless, nodes have to be *marked* to guarantee termination. For this reason, *NextBestOnce marks* nodes if they have no neighbor that is not *marked* and closer to the target t. Since neighbors farer from t than the current node cannot be marked, a node only contacts neighbors that present an improvement or the minimal decline of all neighbors, not only the ones that have not yet seen the message.

Algorithm 1 NextBestOnce(Node p, ID t, Node v, Set B)

```
1: # p predecessor, t target, v current, $B$ marked nodes
2: # $N_v$: neighbors of v
3: if id(v) == t then
4:     routing successful; terminate
5: end if
6: if v.predecessor == null then
7:     v.predecessor = p;
8: end if
9: $S = \{u \in N_v :!B.contains(u)\}$
10: if $S$ NOT EMPTY then
11:     nextNode = $argmin_{u\in S} dist(u,t)$
12:     if $dist(nextNode,t) \geq dist(v,t)$ then
13:         B.add(v)
14:     end if
15: else
16:     B.add(v)
17:     nextNode = v.predecessor; // backtracking
18: end if
19: if nextNode != null then
20:     NextBestOnce(v, t, nextNode, B)
21: else
22:     routing failed; terminate
23: end if
```

In the context of our model, every node has a neighbor within distance C, resulting in a maximal increase of C in distance to the target per step, hence avoiding the large setbacks of D^2-*DFS*.

Indeed, *NextBestOnce* achieves polylog maximal expected routing length, more precisely the maximal expected number of hops is $\mathcal{O}(\log^{\alpha-1} n \log\log n)$. The proof is similar to the one presented in [7], and omitted due to space constraints.

NextBestOnce, described in Algorithm 1, takes as input the predecessor p of the current node, the identifier of the target node t, the current node v, and a set B of *marked* nodes. In each non-terminal step of the algorithm, there are basically two possibilities: The node forwards the message to the neighbor closest to the destination that is not yet marked (ll. 9-14). If this closest neighbor actually is not closer than the current node, the node adds its identifier to B (l. 13, l. 16 respectively if all neighbors are contained in B). It subsequently is not selected as next hop on the path again, unless the message is backtracked. Only during backtracking, other nodes than those closer to the destination and the neighbor with the least increase in distance can be contacted. Backtracking happens in case a node only has neighbors closer to the destination, because otherwise there is an unmarked neighbor, by the condition that nodes are *marked* only after their neighbors

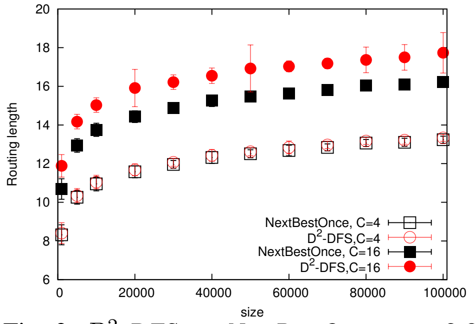

Fig. 2: D^2-*DFS* vs. *NextBestOnce*, $\alpha = 2.3$

closer to t have been *marked*.

The routing fails if the current node is the initiator and all of its neighbors have been contacted (l. 22, this only happens if t does not exist in the connected component).

V. Simulations

Although Sections III and IV give an asymptotic analysis of the routing algorithm, it is unclear how this asymptotic bounds relate to the performance on graphs of a realistic size. We hence perform simulations compare *NextBestOnce* and D^2-*DFS* for realistic network sizes.

Setup: The simulations were performed using GTNA [9], all code is available online[1]. The graphs were generated as follows: For each α, and n, a graph consisting of the nodes with respective long-range links was generated first. In a second step, one graph with short-range links was generated for each value of C. The routing algorithms then were evaluated by creating routing requests to 5 randomly chosen destinations for each node, so $5n$ source-destination-pairs were taken as a sample.

C was chosen to be 1 to 10, 16 and 32, α between 2.1 and 2.5, in steps of 0.05, and n was varied between 1k and 100k. In real-world social networks a value of α between 2.2 and 2.3 has been observed, hence these values are preferably chosen for exemplary evaluation. Please note that α is an artefact of the social graph and cannot be altered in the system design. The results were averaged over 30 to 100 runs.

Results: Indeed, *NextBestOnce* has a lower average routing length than the original D^2-*DFS* for all considered settings in our simulations. Figure 2 displays this performance for network sizes between 1000 and 100000, using $\alpha = 2.3$ and $C = 4, 16$. The performance is very similar for $C = 4$. Nevertheless, *NextBestOnce* has a slightly lower average routing length for all considered network sizes. In case of $C = 16$, the difference between the algorithms is clearly noticeable, with *NextBestOnce* performing over 10% better than D^2-*DFS*. Because the standard deviation of D^2-*DFS* is generally higher than for *NextBestOnce* at $C = 16$, the presented results for D^2-*DFS* are averaged over 100 runs instead of 30. Remarkably, there still are some cases in which the standard deviation is extremely high, indicating several incidents of adverse node placements as described in Section III, which cause the average routing length to increase drastically (n.b. $n \in 20k, 50k, 100k$). This happens only in a small number of runs, so the probability of such a situation to happen at smaller network sizes is low. Considering each single node, it remains constant and hence overall is increasing with the network size. Nevertheless, the average maximal number of routing steps over 100 runs increases at least linearly with the network size, from about 270 steps for 10k and $C = 16$ to more than 10,000 steps for 100k. This shows clearly that such unbeneficial scenarios exist.

VI. Conclusion

This paper deals with routing on connection restricted topologies, especially Darknets. This represents a difficult problem, due to the restriction to establish connections solely between nodes if the respective owners share a trust relationship in real life. The paper introduces a new formal model, which extends the small-world model of Kleinberg to better reflect the realistic properties of Darknets. Both D^2-*DFS*, the routing of Freenet, which is the only currently deployed Darknet, and the newly proposed algorithm *NextBestOnce* are analyzed in the context of the model. The complexity analysis shows that while *NextBestOnce* has an expected polylog routing length, D^2-*DFS* is unable to achieve this performance asymptotically. A simulation study exhibits the polylog performance of *NextBestOnce* The simulations additionally show that situations exist that are highly adverse for D^2-*DFS*. This leads us to the conclusion that *NextBestOnce* is the better choice if either guaranteed polylog routing length are required, or the systems may grow to large network sizes.

In summary, we are positive that the new model will prove to be a useful asset for future analyses of routing protocols on connection restricted topologies, and that *NextBestOnce* represents a promising intermediate step towards enhancing routing in such networks.

References

[1] J. Kleinberg, "The small-world phenomenon: An algorithmic perspective," in *Symposium on Theory of Computing*, 2000.

[2] B. C. Popescu, B. Crispo, and A. S. Tanenbaum, "Safe and private data sharing with turtle: Friends team-up and beat the system," in *Workshop on Security Protocols*, 2004.

[3] T. Isdal *et al.*, "Privacy-preserving p2p data sharing with oneswarm," in *SIGCOMM*, 2010.

[4] E. Vasserman *et al.*, "Membership-concealing overlay networks," in *CCS*, 2009.

[5] P. Mittal, M. Caesar, and N. Borisov, "X-vine: Secure and pseudonymous routing using social networks," *CoRR*, 2011.

[6] I. Clarke *et al.*, "Private communication through a network of trusted connections: The dark freenet," http://freenetproject.org/papers.html,10-12-2010.

[7] S. Roos and T. Strufe, "Provable polylog routing for darknets," in *HotPOST*, 2012.

[8] S. Roos, "Analysis of routing on sparse small-world topologies," Master's thesis, Technische Universität Darmstadt, 2011.

[9] B. Schiller *et al.*, "GTNA: A Framework for the Graph-theoretic Network Analysis," in *Springsim*, 2010.

[1] http://www.p2p.tu-darmstadt.de/research/gtna/