# Always be Two Steps Ahead of Your Enemy* <br> Maintaining a routable overlay under massive churn in networks with an almost up-to-date adversary 

Thorsten Götte ${ }^{\dagger} \quad$ Vipin Ravindran Vijayalakshmi ${ }^{\ddagger}$ Christian Scheideler ${ }^{\text {§ }}$


#### Abstract

We investigate the maintenance of overlay networks under massive churn, where an adversary may churn a constant fraction $\alpha n$ of nodes over the course of $O(\log n)$ rounds. In particular, the adversary has an almost up-to-date information of the network topology as it can observe only a slightly outdated topology that is at least 2 rounds old with a provably minimal restriction that new nodes can only join the network via nodes that have taken part in the network for at least one round. We show that it is impossible to maintain a connected topology if adversary has up-to-date information about the nodes' connections. As our main result we present an algorithm that constructs a new overlay- completely independent of all previous overlays - every 2 rounds. Furthermore, each node sends and receives only $O\left(\log ^{3} n\right)$ messages each round. As part of our solution we propose the Linearized DeBruijn Swarm (LDS), a highly churn resistant overlay, which will be maintained by the algorithm. However, our approaches can be transferred to a variety of classical P2P topologies where nodes are mapped into the $[0,1)$-interval.


## 1 Introduction

Peer-to-peer (P2P) networking has proven to be a useful technique to construct resilient decentralized systems. In a P2P architecture the nodes are connected via the Internet and form a logical network topology, also known as an overlay network. Within the overlay each node has a logical address and logical links that allows it to search and store information in the network.

A key requirement for all applications that rely on P2P networks is reliable communication between all nodes, i.e., each node should be able to send a message to another node at all times. This is complicated by the fact that in every large-scale system, errors and attacks are the rule rather than the exception. At the same time there is usually no or only little admission control for new participants. This implies a massive amount of churn, i.e., nodes joining and leaving the network at any given time. In fact, empirical studies have shown that $50 \%$ of all nodes are subjected to churn over the course of an hour [20]. This alludes for robust and distributed protocols that maintain connected overlays in spite of heavy churn.

In this work we deal with the problem of maintaining a routable overlay under adversarial churn. We define an overlay as routable, if each node in every round is able to send a message to a given logical address $p \in[0,1)$. In each round the adversary picks a set of nodes that leave the network and proposes a set nodes that join the network.

It is easy to see that an adversary that knows all connections between the nodes can simply partition the network by churning out the neighborhood of a node. Previous literature, e.g., 3, 4, 6], therefore

[^0]considered models with an additional restriction, where the adversary has slightly outdated information about the nodes' connections. In particular, the adversary could access all information that is at least $O(\log \log n)$ rounds old, where $n \in \mathbb{Z}^{+}$is the minimal number of nodes in the network in each round. This includes the nodes' connections, internal states, random decisions, the contents of all messages, etc. Within these $O(\log \log n)$ rounds the nodes execute a distributed algorithm that completely rearranges the network topology. However, we remark that the techniques presented in [3, 2, 6] cannot be used if one wants to grant the adversary access to even more recent information.

To overcome above mentioned restriction, we propose a trade-off in the form of a $(a, b)$-late omniscient adversary that has an almost up-to-date information about the network topology, but is more outdated with regard to all other aspects. In particular, it has full knowledge of the topology after $a$ rounds and complete knowledge of messages, internal states, etc. after $b$ rounds. In the real world, an adversary with similar properties could, e.g., be an agency eavesdropping on an Internet exchange points. They can see who communicates based on the involved IP-addresses but are unable to decrypt the messages (or take longer to decrypt them).

Our main contribution is a distributed overlay maintenance algorithm that completely rearranges the network every 2 rounds and can therefore handle a $(2, O(\log n))$-late adversary. Furthermore, the algorithm allows for routing a message to a logical address $p \in[0,1)$ within $O(\log n)$ rounds. The algorithms are randomized and the results hold with high probability (w.h.p. ${ }^{1}$. The overlay we consider in this work is an extension of the Linearized DeBruijn Graph presented in Richa et al. [17] (which by itself is based on the DeBruijn Graph and draws ideas from Naor and Wieder [16]) that uses quorums of logarithmic size to send and receive messages. The latter is adapted from Fiat et al. 9, where the authors use this approach for the Chord-Overlay. We further present a robust algorithm that minimizes the number of messages sent in every step. Our approach uses several structural properties of the overlay as well as a careful analysis of non-independent events to ensure fast reconfiguration of the network.

### 1.1 Model

We assume that time proceeds in synchronous $\mathbb{2}^{2}$ rounds and observe a dynamic set of nodes $\mathcal{V}:=$ $\left(V_{0}, V_{1}, \ldots\right)$ such that $V_{t}$ is the set of nodes in round $t$. Each node is identified by a unique and immutable identifier denoted by $I D$. A node $u \in V_{t}$ can send a message to a node $v \in V_{t}$ only if it knows the $I D$ of node $v$. In a real-world network these $I D$ s could, e.g., be the nodes' IP addresses. This results in series of graphs $\mathcal{G}:=\left(G_{0}, G_{1}, \ldots\right)$ with $G_{i}=\left(V_{i}, E_{i}\right)$ and $E_{t}:=\{(u, v) \mid u$ sends a message to $v$ in round $t\}$. Observe that each $G_{i}$ is a directed graph. Creating an edge may be compared to sending a UDP (User Datagram Protocol) message to the desired receiver or establishing a TCP connection. We assume that a node can create edges to $O(\log n)$ different nodes in each round and can send $O$ (polylog$n)$ bits via each edge. Note that throughout this paper we assume that an $I D$ is of $\operatorname{size} O(\log n)$.

Our model assumes that the set $\mathcal{V}$ is determined by an adversary. This implies, in every round $t$ the adversary can propose a set $O_{t} \subset V_{t-1}$ that leaves the network and a set $J_{t} \subset V_{t}$ that joins the network, i.e., $V_{t}:=V_{t-1} \backslash O_{t} \cup J_{t}$. In particular, the adversary has to comply to the following rules.

## 1. Lateness.

As mentioned in the introduction, we consider the adversary to be $(2, O(\log n))$-late omniscient, i.e., the adversary has slightly outdated knowledge of the topology, i.e., the series of graphs $\mathcal{G}:=\left(G_{0}, G_{1}, \ldots\right)$ created through the communication between nodes. In particular, since our adversary is 2-late in round $t$ the adversary has full knowledge of all graphs until $G_{t-2}$. Further, it has no knowledge of the nodes' internal states and the contents of messages for $O(\log n)$ rounds, i.e., the adversary learns the content of message sent in round $t$ only in round $t+O(\log n)$.

[^1]
## 2. Churn Rate.

For all $V_{t} \in \mathcal{V}$ it holds that $\left|V_{t}\right| \in[n, \kappa n]$, where $\kappa \geq 1$ is a small constant. In other words, the number of nodes stays within $\Theta(n)$. For a suitable value $T \in O(\log n)$, we assume that $V_{t+T} \cap V_{t} \geq(1-\alpha) n$, where $\alpha \in[0,1)$ is a fixed constant. This allows the churn to be $O(n)$ in each round as long as there is a stable set of size $\Theta(n)$ that remains in the network for at least $T$ rounds.

## 3. Bootstrap Phase.

We assume that until a round $B \in O\left(\log ^{2} n\right)$ the adversary is inactive and no churn occurs, also known as the bootstrap phase. We would like to remark that several other works also assume a bootstrap phase to prepare the random sampling (cf. [3, 2, 6]). Only after the conclusion of the bootstrap phase, is the adversary allowed to begin churning nodes in or out of the network.

## 4. Restricted Join.

Additionally, we assume that a new node $v \in V_{t} \backslash V_{t-1}$ can only join the network via a node $w \in V_{t} \cap V_{t-2}$, i.e., the node $v$ joins via node $w$ in round $t$. In Section 2 we show that for our purposes this is a necessary condition. Finally, the number of nodes that join the network via the same node $v \in V_{t}$ is a constant $\phi \in O(1)$.

We remark that our model incorporates observations from Stutzbach and Reza [20] that new nodes join and leave very frequently but there is a (relatively) stable set of older nodes. Note that to the best of our knowledge this is a significantly more flexible model compared to other related work. Given the above mentioned constraints, a round $t$ consists of the following four steps.

1. At the beginning of each round the adversary can select a set of nodes $O_{t} \subset V_{t-1}$ that leave the network in round $t$. These nodes do not receive any messages and leave the network immediately. Further, the adversary may propose a set of nodes $J_{t}$ that joins the network in round $t$. For each node $v \in J_{t}$ the adversary selects a bootstrap node $w \in V_{t} \backslash J_{t}$ (satisfying the necessary conditions for Restricted Join) that receives a reference to $v$.
2. Next, all nodes that are still in the system receive all messages sent in the previous round. Note that this even holds for messages that were sent by nodes that were churned out in the current round. In other words, if a node can send out a message, it will be received.
3. After receiving all messages, a node can perform calculations on its local variables and the received messages.
4. Finally, each node may send messages to other nodes. Recall that sending a message to another node implicitly creates an edge in the graph $G_{t+1}$. Every message sent in round $t$ is received in the round $t+1$. Furthermore, due to the lateness condition these edges can only be seen by the adversary at the beginning of the round $t+3$.

### 1.2 Related Work

There has been extensive work on analyzing overlay networks under high adversarial churn. As already mentioned in the introduction, these works had a variety of different model assumptions. See Augustine and Sivasubramaniam [3 for a comprehensive survey on previous results. In the following, we only concentrate on models closely related to ours.

First, there was a series of papers (cf. [19, 9, 4]) that assumed only a subset of nodes is subjected to adversarial churn. However, these nodes could also act byzantine and try to sabotage the overlay's maintenance and the routing by sending corrupted messages. A general assumption was that up to a

[^2]Table 1: Overview of different models in the literature

| Paper | Lateness $^{\mathbf{a}}$ | Churn Rate $^{\mathbf{b}}$ | Immediate $^{\text {a }}$ |
| :---: | :---: | :---: | :---: |
| $[3$ | $(O(\log \log n), O(\log \log n))$ | $(\alpha n, O(\log \log n))$ | Yes |
| 6 | $(O(\log \log n), O(\log \log n))$ | $\left(n-\frac{n}{\log n}, O(\log \log n)\right)$ | $\mathrm{No}^{\mathrm{c}}$ |
| $[1]$ | $(O(\log n), O(\log n))$ | $\left(O\left(\frac{n}{\log n}\right), O(\log n)\right)$ | Yes |
| Our work | $(2, O(\log n))$ | $(\alpha n, O(\log n))$ | Yes |

${ }^{\text {a }}$ An adversary is ( $a, b$ )-late if it has full knowledge of the topology after $a$ rounds and complete knowledge after $b$ rounds.
${ }^{\mathrm{b}}$ The churn rate is $(C, T)$ if the adversary can perform $C$ join/leaves in $T$ rounds.
${ }^{\text {c }}$ Nodes remain in the network for additional $O(\log \log n)$ rounds.
constant fraction of nodes would be malicious. Scheideler [19] present a protocol that spreads these nodes over the network such that each connected subset of logarithmic size contains a constant fraction of non-byzantine nodes. Fiat et al. [9] build upon this work and present a full overlay maintenance algorithm that provided a robust Distributed Hash Table. In their approach, each virtual address $p \in[0,1)$ is maintained by a committee of $O(\log n)$ nodes. We will reuse this idea in our work.

In more recent works all of the nodes are subjected to adversarial churn and not only a fixed set. However, these works usually do not consider byzantine behavior. The adversary in these papers can be described by three properties: The lateness, the churn rate, and if it is immediate. We say adversary is $(a, b)$-late if it has full knowledge of the topology after $a$ rounds and complete knowledge of all sent messages, internal states, etc. after $b$ rounds. The churn rate is $(C, T)$ if the adversary can perform $C$ join/leaves in $T$ rounds. Last, an adversary is immediate if churned out nodes have to leave the network immediately and without the possibility to send and receive more messages. Table 1 shows an overview over the different models. Note that the table is only for comparison as it simplifies some of the models and does not depict all of their respective nuances. However, these simplifications do not weaken the adversary.

Augustine et al. [2] present an algorithm that builds and maintains an overlay in the presence of a nearly completely oblivious adversary. Here, the overlay no longer has a fixed structure but is an unstructured expander graph of constant degree. Note that this overlay has no virtual addressing. However, in [1] the authors present a scheme that allows to quickly search for data in these networks.

Further, Drees et al. [6] build a structured expander, a so-called $H_{d}$-Graph, which is the union of $d$ random rings. Their adversary is not only $O(\log \log n)$-late with regard to communication, it also has access to all nodes' memory and all sent messages after $O(\log \log n)$. Nodes that are churned out in round $t$ may remain in the network until some round $T \in O(t+\log \log n)$. Thus, it is not immediate.

Last, the SPARTAN framework presented in [3] probably bears the greatest resemblance with our work. In SPARTAN the nodes maintain a logical overlay resembling a butterfly network. To ensure robustness each of the butterfly's virtual nodes is simulated by $O(\log n)$ nodes. The key difference between our work and SPARTAN is the adversary's lateness. Similar to [6], the SPARTAN framework assumes the adversary to be $(O(\log \log n), O(\log \log n))$-late, but in return allows the churn to be as high as $\alpha n$ in $O(\log \log n)$ rounds. However, unlike [6], SPARTAN allows the adversary to be immediate.

### 1.3 Our Contribution

In this work we present an algorithm which given a dynamic set of nodes $\mathcal{V}:=\left(V_{0}, V_{1}, \ldots\right)$ chosen by a $(2, O(\log n))$-late adversary, creates a series of graphs $\mathcal{G}:=\left(G_{0}, G_{1}, \ldots\right)$ with $G_{i}:=\left(V_{i}, E_{i}\right)$, such that it holds w.h.p. that $G_{i}$ is routable, i.e., each node can send a message to a logical address $p \in[0,1)$.

The paper is organized as follows.

- In Section 1.4 we introduce the Linearized DeBruijn Swarm (LDS). This graph topology is based
on the linearized DeBruijn Graph presented by Richa et al. [17] and the concept of swarms used by Fiat et al. [9] for the Chord overlay network.
- In Section 2 we show that our model assumptions are necessary in order the solve the problem. In particular, we show that any adversary can partition a network where nodes can join via nodes that themselves just joined one round ago. Further, we prove that our model requires the adversary to be at least 1-late with regard to the topology.
- In Section 3 we present a routing algorithm for the LDS, which optimizes the congestion if we want guaranteed message delivery. In this section, we also define when a dynamic overlay is routable.
- Section 4 we present our main contribution, an algorithm that rearranges the graph topology such that it is completely rebuilt every 2 rounds but still allows routing. The message complexity is $O\left(\log ^{3} n\right)$ messages per node and round w.h. $x^{4}$.


### 1.4 Definitions and Preliminaries

In this section we present some definitions and results from probability theory that we will use in the analysis of our algorithms. During our analysis we deal with both dependent and independent random variables. The two following general classes of random variables will prove to be useful. First, there is Negative Correlation:

Definition 1 (Negative Correlation, see e.g. [18, p. 31]) A set of random variables $X_{1}, X_{2}, \ldots, X_{n}$ are said to negatively correlated if any subset $X_{S} \subseteq X$ it holds that,

$$
\mathbf{E}\left[\prod_{i \in S} X_{i}\right] \leq \prod_{i \in S} \mathbf{E}\left[X_{i}\right]
$$

Here, $S \subseteq[1, n]$ is the set of indices in $X_{S}$.
Further, there is the slightly stronger notion of Negative Association:
Definition 2 (Negative Association [12, 21]) A set of random variables $X_{1}, X_{2}, \ldots, X_{n}$ are said to negatively associated (NA) if for any two functions $f, g$ both monotonically increasing (or both monotonically decreasing) defined on disjoint subsets of $X$ it holds that,

$$
\mathbf{E}[f(X) \cdot g(X)] \leq \mathbf{E}[f(X)] \cdot \mathbf{E}[g(X)]
$$

Note that all independent and hyper-geometric random variables are always NA [7].
Corollary 1 (NA implies Negative Correlation [21]) Let $X_{1}, X_{2}, \ldots, X_{n}$ be a set of $N A$ random variables. Then for all $X_{S} \subset X$ such that $i \neq j$ it holds that,

$$
\mathbf{E}\left[\prod_{i \in S} X_{i}\right] \leq \prod_{i \in S} \mathbf{E}\left[X_{i}\right]
$$

The following propositions from Joag-Dev and Proschan [12] will be extensively used in many of our proofs.

Proposition $2([\mathbf{1 2}, ~ 7])$ If $X:=\left(X_{1}, \ldots, X_{n}\right)$ and $Y:=\left(Y_{1}, \ldots, Y_{m}\right)$ are negatively associated sets of random variables that are mutually independent, then the vector $(X, Y):=\left(X_{1}, \ldots, X_{n}, Y_{1}, \ldots, Y_{n}\right)$ are also negatively associated.

[^3]

Figure 1: A node $v$ is connected each node in red areas. Note that the swarms $S(v), S\left(\frac{v}{2}\right)$, and $S\left(\frac{v+1}{2}\right)$ are real subsets. These sets are chosen such that they contain $O(\log n)$ nodes w.h.p. Recall that in a classical Debruijn Graph it would only be connected to the nodes left and right of $v$ i.e., $\frac{v}{2}$ and $\frac{v+1}{2}$.
Proposition $3([\mathbf{1 2}, \mathbf{7}])$ Let $\left(X_{1}, \ldots, X_{n}\right)$ be negatively associated random variables. For some $k \leq n$, let $I_{1}, \ldots, I_{k} \subseteq[n]$ be disjoint index sets. For $j \in[k]$, let $f_{j}: \mathbb{R}^{\left|I_{k}\right|} \mapsto \mathbb{R}$ be functions that are all nondecreasing or all non-increasing. Define $Y_{j}:=f_{j}\left(X_{i}: i \in I_{j}\right)$. Then the random variables $\left(Y_{1}, \ldots, Y_{k}\right)$ are negatively associated.

Further, we make use of the following Chernoff Bounds, which are defined as follows.
Lemma 4 (Chernoff-Hoeffding Bounds [7, 15]) Let $X:=\sum X_{i}$ be the sum of negatively correlated random variables with $X_{i} \in\{0,1\}$. Then it holds that for any $0<\delta<1$,

$$
\operatorname{Pr}[X \geq(1+\delta) \mathbf{E}[X]] \leq e^{-\frac{\delta^{2} \mathbf{E}[X]}{2}}
$$

and

$$
\operatorname{Pr}[X \leq(1-\delta) \mathbf{E}[X]] \leq e^{-\frac{\delta^{2} \mathbf{E}[X]}{3}}
$$

and for any $\delta \geq 1$, it holds:

$$
\operatorname{Pr}[X \geq(1+\delta) \mathbf{E}[X]] \leq e^{-\frac{\delta \mathbf{E}[X]}{3}}
$$

Throughout this work we assume that each node in the network is aware of $n$ and $\kappa$, i.e., the lower and upper bound on the number of nodes currently in the network. We make this simplification due to Stutzbach and Reza [20], that the number of nodes stays relatively stable. Furthermore, in order to simplify notations, we define $\lambda:=2 \log \kappa n 5$. We would like to remark that all of our algorithms presented in this manuscript may be adapted to work with close estimates of $\lambda$ and $\frac{\lambda}{n}$ using approaches presented in [17, 9, 14, 13].

## DeBruijn Swarm

We now present our overlay, the Linearized DeBruijn Swarm (LDS), which is a combination of a wellanalyzed network overlay of low degree, i.e., the Linearized DeBruijn Graph (LDG) presented in [17, 8] with techniques from robust overlays, i.e., the usage of logarithmic quorums that simulate a single node [9. Note that the LDG is inspired by but not equivalent to the classical DeBruijn Graph. The notion of swarms were also described in Fiat et al. 9].

In the remainder of this section we present the LDS's topology and show some of its basic properties. Each node $v \in V$ chooses a position $p_{v} \in[0,1)$ uniformly and independently at random. Note that for the sake of convenience, the position of a node $v \in V$, we just write $v$ instead of $p_{v}$. It should always be clear from the context if we refer to the node, its ID, or its (current) position. However, we will attempt to make the context sufficiently clear in order to reduce the level of ambiguity.

All nodes can calculate the distance to another node via the distance function $d: V^{2} \rightarrow[0,1)$. Given two nodes $v, w \in V$ the distance function $d$ returns the shortest distance (hop counts) between $v$ and $w$ in the $[0,1)$-torus.

[^4]Formally, the function $d: V \times V \mapsto[0,1)$ is defined as follows,

$$
d(v, w):= \begin{cases}|v-w| & \text { if }|v-w| \leq 1 / 2  \tag{1}\\ 1-|v-w| & \text { otherwise }\end{cases}
$$

Furthermore, the distance function also satisfies the triangular inequality, i.e.,

$$
\begin{equation*}
d(v, w) \leq d(v, z)+d(z, w) \quad \text { for all } z \in V \tag{2}
\end{equation*}
$$

For convenience we introduce the following notions for the relation between two nodes $u, v \in V$. If $|u-v| \leq 1 / 2$, then $u$ is left (clockwise) of $v$ if $u<v$ and right (clockwise) otherwise. For $|u-v|>1 / 2$ the relation is reversed. Further, the set $\langle u, v\rangle \subset V$ contains all nodes which are right of $u$ and also left of $v$. Given a node $w$ we say that $u$ is closer to $w$ than $v$ if $d(u, w)<d(v, w)$. Last, we call a node $u$ the closest neighbor of $v$ if there exists no other node $u^{\prime} \in V$ closer to $v$ than $u$, i.e., $u:=\underset{k \in V}{\operatorname{argmin}} d(k, v)$.

In the LDG presented by Richa et al. [17, each node $v$ connects to exactly six other nodes. Namely, the two closest nodes left and right of $p_{v}$ and the two closest node left and right of the points $\frac{p_{v}}{2}$ and $\frac{p_{v}+1}{2}$ respectively. We extend this structure such that, each node connects to $O(\log n)$ closest neighbors. For a given point $p \in[0,1)$ we call $S(p) \subset V$ the swarm of $p$. It holds that $v \in S(p)$ if and only if $d(v, p) \leq \frac{c \lambda}{n}$. Here, $c>1$ is a robustness parameter which should be chosen as small as possible. These swarms (and not the nodes) will be the building blocks of our overlay. We call the swarms $S(p)$ adjacent to $S\left(p^{\prime}\right)$ if there is an edge $(v, w)$ between every node $v \in S(p)$ and $w \in S\left(p^{\prime}\right)$. Note that each swarm $S(p)$ spans an interval of length $\frac{2 c \lambda}{n}$ as it consists of two intervals of length $\frac{c \lambda}{n}$ to the left and right of $p$ respectively. Sometimes it will be necessary to distinguish between the nodes in the left and right interval of $p$, so we define $S^{L}(p):=\{v \in S(p) \mid v$ is left of $p\}$ and $S^{R}(p):=\{v \in S(p) \mid v$ is right of $p\}$ as the left and right side of $S(p)$. Given this notion of swarms, we can now formally define the adapted overlay. Formally the LDS is defined as follows:

Definition 3 (Linearized DeBruijn Swarm) Let $V \subset[0,1)$ be a set of points with $|V|=n$ and $\lambda:=\log n$. Then, the $L D S G_{c}:=\left(V, E_{L} \cup E_{D B}\right)$ with parameter $c \in \mathbb{N}$ has the following properties:

- $(v, w) \in E_{L} \Longleftrightarrow w \in V$ and $d(v, w) \leq \frac{2 c \lambda}{n}$.
- $(v, w) \in E_{D B} \Longleftrightarrow w \in V$ and $d\left(\frac{v+i}{2}, w\right) \leq \frac{3 c \lambda}{2 n}$ with $i \in\{0,1\}$.

A Linearized DeBruijn Swarm is illustrated in Figure 1. Over the course of this paper we will refer to the edges in $E_{L}$ as list edges, whereas the edges in $E_{D B}$ as DeBruijn edges.

From Definition 3 we state the following lemma.
Lemma 5 (Swarm Property) Consider any point $p \in[0,1)$ and its swarm $S(p) \subset V$. Then $S(p)$ is adjacent to $S\left(\frac{p}{2}\right)$ and $S\left(\frac{p+1}{2}\right)$.

Proof: Let $p \in[0,1)$ be any point and $v \in S(p)$ be a node in $p$ 's swarm. From Definition 3 this implies that $d(p, v) \leq \frac{c \lambda}{n}$. We now show that node $v$ has a connection to every node in $S\left(\frac{p}{2}\right)$ and $S\left(\frac{p+1}{2}\right)$. For the proof we only analyze adjacency to $S\left(\frac{p}{2}\right)$, since the other case is analogous. We distinguish between the following two scenarios.

1. If $|p-v| \leq \frac{1}{2}$, then from (1) it holds that,

$$
\begin{equation*}
d\left(\frac{p}{2}, \frac{v}{2}\right):=\left|\frac{p}{2}-\frac{v}{2}\right|=\frac{1}{2}|p-v| \leq \frac{1}{2} \frac{c \lambda}{n} . \tag{3}
\end{equation*}
$$

Let $u$ be any node in $S\left(\frac{p}{2}\right)$, then $d\left(u, \frac{p}{2}\right) \leq \frac{c \lambda}{n}$. Then using (2) and (3) we get,

$$
d\left(u, \frac{v}{2}\right) \leq d\left(u, \frac{p}{2}\right)+d\left(\frac{p}{2}, \frac{v}{2}\right)
$$

$$
\begin{align*}
& \leq \frac{c \lambda}{n}+\frac{c \lambda}{2 n} \\
& =\frac{3 c \lambda}{2 n} \tag{4}
\end{align*}
$$

From Definition 3, since node $v$ has a DeBruijn edge to each node $w \in V$ with $d(v / 2, w) \leq 3 c \lambda / 2 n$, then using (4) the lemma follows.
2. Otherwise, $|p-v|>\frac{1}{2}$.

Observe that this only occurs if either $p \in\left[0, \frac{c \lambda}{n}\right]$ or $p \in\left[1-\frac{c \lambda}{n}, 1\right)$, i.e., the point $p$ is close to 0 or 1 and $v$ lies on the opposite site of the interval. We distinguish between the following two cases.
(a) If $p \in\left[0, \frac{c \lambda}{n}\right]$, then it also holds that $\frac{p}{2} \in[0, p]$. Implies,

$$
\begin{equation*}
d\left(v, \frac{p}{2}\right) \leq d(v, p) \leq \frac{c \lambda}{n} \tag{5}
\end{equation*}
$$

Then, by the triangle inequality and inequality (5), it holds that for every node $u \in S\left(\frac{p}{2}\right)$,

$$
d(u, v) \leq d\left(u, \frac{p}{2}\right)+d\left(\frac{p}{2}, v\right) \leq \frac{2 c \lambda}{n} .
$$

From Definition 3, the node $u$ is then a list neighbor of $v$ and the lemma follows.
(b) Otherwise, if $p \in\left[1-\frac{c \lambda}{n}, 1\right)$ then it holds that $\frac{p}{2} \in\left[\frac{1}{2}-\frac{c \lambda}{2 n}, \frac{1}{2}\right)$. Now consider the distance between $\frac{p}{2}$ and $\frac{v+1}{2}$. Here, it holds

$$
d\left(\frac{v+1}{2}, \frac{p}{2}\right):=\left|\frac{v+1}{2}-\frac{p}{2}\right|=\frac{1}{2}|v+1-p| .
$$

Observe that since $v<p$ and $p<v+1$, it holds that $|(1+v)-p|$ is equivalent to $1-|v-p|$. Therefore, the inequality simplifies to

$$
d\left(\frac{v+1}{2}, \frac{p}{2}\right)=\frac{1}{2}(1-|v-p|)=\frac{1}{2} d(v, p) .
$$

Since $\frac{1}{2} d(v, p) \leq \frac{c \lambda}{2 n}$, applying the triangle inequality we get that for every node $u \in S(p / 2)$,

$$
\begin{align*}
d\left(u, \frac{v+1}{2}\right) & \leq d\left(u, \frac{p}{2}\right)+d\left(\frac{v+1}{2}, \frac{p}{2}\right) \\
& \leq \frac{c \lambda}{n}+\frac{c \lambda}{2 n} \\
& =\frac{3 c \lambda}{2 n} \tag{6}
\end{align*}
$$

From Definition 3 since node $v$ has a DeBruijn edge to each node $w \in V$ with $d(v+1 / 2, w) \leq$ $3 c \lambda / 2 n$, then using (6) the lemma follows.

The next lemma shows that if nodes are assigned to the [ 0,1 )-interval uniformly and independently at random, then all swarms have roughly the same size w.h.p.

Lemma 6 (Swarm Size) Let the swarm size to be $\frac{c \lambda}{n}$ with $c \geq 12 k$ and assume all nodes pick their positions uniformly and independently at random. Then for any point $p \in[0,1)$ it holds that,

$$
\operatorname{Pr}\left[\frac{1}{2} c \lambda<|S(p)|<2 c \lambda\right] \geq 1-\frac{2}{n^{k}},
$$

where $|S(p)|$ denoted the number nodes in $S(p)$.


Figure 2: Example for the first four steps of a trajectory

Proof: The proof follows from a standard application of the Chernoff bound. Given that each node picks its position uniformly and independently at random in $[0,1)$-interval, the probability that a node chooses a point in an interval of length $\frac{c \lambda}{n}$ is exactly $\frac{c \lambda}{n}$ 6.

Let $X$ be a random variable that counts the number of nodes in $S(p)$. For each $v \in V$, let $X_{v}$ be a binary random variable such that, $X_{v}=1$, if $v$ picked a position in a swarm $S(p)$ and 0 , otherwise. Clearly, it holds that $X:=\sum_{v \in V} X_{v}$ and $\mathbf{E}[X]=c \lambda$ for every $p$. Furthermore, it holds that $X:=$ $\sum_{v \in V} X_{v}$ is the sum of independent random variables. Hence, using Lemma4 it holds that for $c \geq 12 k$,

$$
\begin{equation*}
\operatorname{Pr}\left[X \leq \frac{1}{2} \mathbf{E}[X]\right] \leq e^{-\frac{\mathbf{E}[X]}{12}} \leq e^{-k \lambda}=n^{-k} \tag{7}
\end{equation*}
$$

Moreover, it holds that for $c \geq 2 k$,

$$
\begin{equation*}
\operatorname{Pr}[X \geq 2 \mathbf{E}[X]] \leq e^{-\frac{\mathbf{E}[X]}{2}} \leq e^{-k \lambda}=n^{-k} \tag{8}
\end{equation*}
$$

Then, using (7), (8) and the fact that the expected number of nodes in an interval of size $\frac{c \lambda}{n}$ is $c \lambda$, the union bound yields the desired result.

## Routing in a Linearized DeBruijn Graph

Our routing algorithm described in Section3is based on the classical LDG's routing algorithm presented in [17, 8. Before we go into the details of our algorithm, we will first recall the classical LDG's routing algorithm. Routing in the LDG works by a bitwise adaption of the target address. Recall that we assume each node knows $\lambda$. Therefore, given any destination $p \in[0,1$ ), a node can calculate the first $\lambda$ bits $\left(p_{1}, \ldots, p_{\lambda}\right)$ of $p$ 's binary representation. Then, starting with the least significant bit $p_{\lambda}$, the node $v$ sends the message to the node closest to $x_{1}:=\frac{v+p_{\lambda}}{2}$. For this, it uses the corresponding DeBruijn edge. After that, the message is sent to the node closest to $x_{2}:=\frac{x_{1}+p_{\lambda-1}}{2}$. This goes on until the first bit $p_{1}$. Finally, as a consequence of Lemma 6 there are w.h.p. only $O(\log n)$ hops over list edges left to $p$.

Definition 4 (Trajectory) Let $v \in V$ be a node and $p \in[0,1)$ be an arbitrary point. Further let $\left(v_{1}, \ldots, v_{\lambda}\right) \in\{0,1\}^{\lambda}$ and $\left(p_{1}, \ldots, p_{\lambda}\right) \in\{0,1\}^{\lambda}$ be the $\lambda$ most significant bits of $v$ and $p$ respectively. Then the trajectory $\tau(v, p):=x_{0}, \ldots, x_{\lambda+1} \in[0,1)^{\lambda+2}$ is a series of points defined as follows.

$$
x_{i}:= \begin{cases}v & i=0 \\ \left(p_{\lambda-i+1}, \ldots, p_{\lambda}, v_{1}, \ldots, v_{\lambda-i}\right) & i \leq \lambda \\ p & i=\lambda+1\end{cases}
$$

For each point $x_{i}$ in the trajectory, forward the message to the node closest to it. Then, forward the message along list edges until it reaches the target.

[^5]
## 2 Impossibility Results and Lower Bounds

In this section, we present two fundamental impossibility results for our model. First, we show that it is impossible to maintain a connected overlay under massive churn and a ( $0, \infty$ )-late adversary. This adversary always has up-to-date information about the topology, but is oblivious of everything else, e.g., the sent messages, the nodes' random decisions, etc. Second, we show the necessity that new nodes can only join via bootstrap nodes that are in the network for at least 2 rounds.

We begin with an auxiliary lemma and show that any adversary with a churn rate $(\alpha n, O(\log n))$ can completely exchange the set of nodes within $O(\log n)$ rounds if $\alpha$ is a constant. Therefore it simply churns out the nodes in chunks of size $\alpha n$.

Lemma 7 Consider any ( $a, b$ )-late adversary that proposes a series of nodes $\mathcal{V}:=V_{0}, V_{1}, \ldots$ such that for all $V_{t} \in \mathcal{V}$ it holds that $\left|V_{t}\right| \in \Theta(n)$ and $\left|V_{t} \cap V_{t+T}\right| \geq(1-\alpha) n$ for $T \in O(\log n)$ and some constant $\alpha \in(0,1)$. Then within $O\left(\alpha^{-1} \log n\right)$ rounds the adversary can churn out all nodes from $V_{0}$.

Proof: For simplification assume that $\alpha:=\frac{1}{\beta}$ with $\beta>1$. and that $\frac{n}{\beta}$ is an integer. If this is not the case, the adversary can add/remove some nodes such that it holds.

Now divide the set $V_{0}$ into $\beta$ disjoint subsets $V_{1}^{1}, \ldots, V_{0}^{\beta}$ of size $\frac{n}{\beta}$ each. The adversary's strategy is as follows: For $i \in 1, \ldots, \beta$ churn out $V_{0}^{i}$ in round $T \cdot i$ and replace it with a set $\tilde{V}^{i}$ of the same size. Thus, the following three statements hold:

1. Recall that $\beta$ is a constant. Thus, after $\beta \cdot T$ rounds all nodes $V_{0}$ have been churned out.
2. For each node that is churned out in round $t$ a new node is churned in. This ensures that the number of nodes in each round is $n$.
3. Within a period of length $T$ only one set of $\operatorname{size} \frac{n}{\beta}=\alpha n$ is churned out. Thus, there always is a subset of size $(1-\alpha) n$ which remains in the system for $T$ rounds.

Thus, the strategy fulfills all requirements and churns out all nodes from $V_{0}$ within $O(\log n)$ rounds. This was to be shown.

We now show the impossibility for a $(0, \infty)$-late adversary. The idea behind this proof is as follows: Consider a node $v \in V_{t}$ joining the network in round $t$ via some node $w \in V_{t}$. Then only $w$ and all nodes $w$ communicates with know $v$. A 0-late adversary can immediately detect and churn out these nodes. Thus, no node in the entire system knows $v$. The result is stated in the following lemma.

Lemma 8 A $(0, \infty)$-late adversary with churn rate $(\alpha n, O(\log n))$ for some $\alpha \in(0,1)$, can disconnect any overlay in $O(\log n)$ rounds.

Proof: Let the execution start at round 0 and let $V_{0}$ be the initial set of nodes with $\left|V_{0}\right|:=n$. Now consider the following strategy:

1. Let a node $v$ join the network in round 0 via any node in $V_{0}$.
2. Further, let a node $w$ join via $v$ in round 2 .

We will show that within $O(\log n)$ rounds, a 0 -late adversary can separate $w$ from the network.
Let $D_{2} \subset V_{0}$ be the set of all nodes that $v$ communicates with in round 2. Note that $\left|D_{2}\right| \in O(\log n)$ because we assume that each node can only communicate with a logarithmic number of other nodes in one round. As $v$ is the only node that knows $w$ in round 2 , it holds that $w$ can only be known by nodes from $D^{\prime}:=D_{2} \cup\{v\}$ in round 3 . On the other hand, $v$ may have sent a set of IDs $D_{w} \subset V_{0}$ to $w$ in round 2 . These are the only nodes that $w$ knows.

Since $D^{\prime}$ is of logarithmic size, there exists an $\alpha \in(0,1)$ such that it is well within the permitted churn size per round. Let all nodes in $D^{\prime}$ be churned out in round 3, i.e., before they can communicate with any more nodes. This ensures that $w$ 's ID cannot be known to any node in the system.

Further, $w$ knows only the IDs received from $v$ as it has only communicated with $w$ until now. Now continue as follows. In each round until all nodes from $V_{0}$ are gone:

1. Churn out every node $w$ communicates with. This ensures that no new node will learn $w$ 's ID.
2. Use the strategy given above to churn out as much nodes from $V_{0}$ as possible and churn in the same amount of new nodes. Note that the total number of nodes does not change.

Using Lemma 7 one can easily verify that within $O(\log n)$ rounds all nodes from $V_{0}$ are gone. Since all ID that $w$ knows belong to nodes from $V_{0}$ and no node in the network knows $w$, it is separated from the network. This concludes the proof.

We continue with the restrictions for the joining nodes. The result is stated in the following lemma.
Lemma 9 Let $v \in V$ be a node that joined in round $t$. Now assume a model where in round $t+1$ a new node $w \in V$ can join the network via $v$. Then $a(\infty, \infty)$-late adversary with churn rate $(\alpha n, O(\log n))$ for some $\alpha \in(0,1)$ can disconnect any overlay after $O(\log n)$ rounds.

Proof: Consider a set of nodes $v_{1}, \ldots, v_{T}, v_{T+1}$ for some $T \in O(\log n)$ that join the network one after another such that $v_{i}$ joins via $v_{i-1}$ in round $i$. Further, for each $v_{i}$ let $D_{i}$ be the set $I D$ s that it initially receives from $v_{i-1}$. Let the execution start at round 0 and let $V_{0}$ be the initial set of nodes with $\left|V_{0}\right|:=n$.

We first show that the adversary can create a situation where a node $v_{T+1}$ that joined the network in round $T+1$ only receives IDs of churned out nodes and thereby disconnecting the network. Consider the following strategy: Let $V^{\prime}:=v_{0}, \ldots, v_{T}$ be a set of nodes such that each $v_{t} \in V^{\prime}$ with $t>1$ joins the network in round $t$ via node $v_{t-1}$. To be precise, this implies that at the beginning of round $t, v_{t-1}$ knows the $I D$ of $v_{t}$ an can send a message to $v_{t}$. Any of these messages will arrive $t+1$. Further, each $v_{t} \in V^{\prime}$ is churned out in round $t+2$, i.e., immediately after the round $v_{t+1}$ joined.

Now let $D_{t} \subset V$ be set of all IDs that $v_{t}$ knows in round $t+1$. We now claim that it holds $D_{t} \subset D_{1} \cup\left\{v_{t-1}\right\}$ for all $t>1$.

We proof the claim via induction:

- For the induction's beginning consider $t=1$ and the corresponding node $v_{1}$. Here, the claim trivially holds as $D_{1} \subseteq D_{1}$.
- For the induction's step consider any $v_{t}$ with $t>1$ and assume that the claim holds for $v_{t-1}$, i.e., it holds that $D_{t-1} \subset D_{1} \cup\left\{v_{0}, \ldots, v_{t-2}\right\}$. Now consider the join of $v_{t}$ in round $t$. Any message that reaches $v_{t}$ in round $t+1$ must be sent in round $t$. However, in this round only $v_{t-1}$ knows $v_{t}$ and therefore only $v_{t-1}$ may share the references with $v_{t}$. Thus, $D_{t}$ can only be a subset of $D_{t-1} \cup\left\{v_{t-1}\right\}$. This proves the claim.

Therefore, a node $v_{T}$ that joins in round $T \in O(\log n)$ only receive references to nodes from $V_{0}$ in round $T+1$ (from $v_{T-1}$ ). Note that any further reference can reach $v_{T}$ only in round $T+2$.

By Lemma 7 the adversary may have churned out all nodes from $V_{0}$ for some $T \in O(\log n)$. Now let a new node $v_{T+1}$ join via $v_{T}$ in round $T+1$. Then $v_{T}$ cannot introduce $v_{T+1}$ to any node currently in the network and further cannot introduce any node $w \in V_{T+1}$ to $v_{T+1}$. Thus, once $v_{T}$ is churned out in round $T+2$, it holds that $v_{T+1}$ is isolated from $V_{T+2}$. This was to be shown.

We would like to remark that this impossibility is different from the similar statement in 2] because we allow a node to communicate with $O(\log n)$ different nodes instead of constantly many.

## 3 Routing and Sampling in the LDS under Churn

In this section, we present a low-congestion routing algorithm Alg-Routing. The algorithm delivers each message w.h.p. even in the presence of churn and a changing communication structure. We also present a sampling algorithm Alg-Sampling that allows each node to send a message to a uniformly picked random node. The underlying technique is adapted from King et al. [14, 13].

Our algorithm must perform routing over a dynamic series of graphs $\mathcal{D}:=\left(D_{1}, D_{2}, \ldots\right)$ where each $D_{i}$ is a LDS. However, there are two problems we need to address, i.e., the churn orchestrated by the adversary and the dynamic reconfiguration of the overlays. The obvious solution would be to send the

## Listing 1: Alg-Routing

```
1Desc: This algorithm is executed on a series of routable graphs \(\mathcal{D}:=\left(D_{1}, H_{1}, \ldots\right)\).
        The algorithm routes a message \(m\) from any node \(u \in V_{1}\) to swarm \(S_{2 \lambda+2}(p)\).
2
3 Note: The following code is executed by each node \(u \in V^{t}\) every round \(t\). W.l.o.g.
        the Forwarding step executed in even rounds and the Handover step is executed
        in odd rounds. The initial step may be executed in odd or even rounds.
        Messages are delivered to their target swarms in even or odd round
        respectively.
4
5 —— Initial step
6 Upon sending a message \(m\) to \(p\)
\(7\left(d_{1}, \ldots, d_{\lambda}\right) \longleftarrow \lambda\) most significant bits of \(p\)
    if \(t\) is even:
        Send \(\left(p, 0,\left(d_{1}, \ldots, d_{\lambda}\right), m, t\right)\) to all \(w \in S_{t}(v)\)
        else:
            Send \(\left(p, 0,\left(d_{1}, \ldots, d_{\lambda}\right), m, t\right)\) to all \(w \in S_{t+1}(v)\)
12
13
4 Upon receiving \(m:=\left(p, k,\left(d_{1}, \ldots, d_{\lambda_{v}}\right), m, t\right)\)
    if \(k \leq \lambda_{v}\) :
        \(x \longleftarrow \frac{v+d_{k}}{2}\)
        \(\left(w_{1}, \ldots, w_{r}\right) \longleftarrow r\) nodes chosen u.i.r from \(S(x)\)
        Forward \(m^{\prime}:=\left(p, k+1,\left(d_{1}, \ldots, d_{\lambda_{v}}\right), \lambda_{v}\right)\) to all \(w_{i}\)
    else if \(t\) is even:
        Deliver \(m\) all nodes \(w \in S_{t+1}(p)\) in next round
    else:
        Deliver \(m\) all nodes \(w \in S_{t}(p)\)
    ,
    ——Handover Step
    Upon switching to \(D_{i+1}\)
    for each message \(m\)
        \(\left(w_{1}, \ldots, w_{C}\right) \longleftarrow r\) nodes chosen u.i.r from \(S_{t+1}(x)\)
28 Forward \(m\) to all \(w_{i}\)
```

message not only to the closest node of each trajectory point but to the whole swarm. However, observe that this trivial adaption to the LDG routing algorithm fails in the presence of churn. Given that any node on a message's trajectory can be churned out, a fraction of routing requests may never reach their destinations. In particular, if the adversary is aware of the topology, it could even churn out the whole swarm for a given trajectory point. Therefore, we introduce the notion of a good swarm adapted from Fiat et al. 9]. In their work, a swarm is good if at least a fixed constant fraction of its nodes take part in the next round and hence refer to such nodes as good.

Here, we need a slightly stronger notion as we require the good nodes to be somewhat well spread over the swarm to enable our fast construction. To be precise, we say the left (right) side of a swarm is good if a constant fraction of its nodes is good and the full swarm is good, if both its left and right side are good. Further, a LDS is good if all its swarms are good. This property implies that there is always at least a constant fraction of good nodes in each swarm that can forward the message.

Besides the churn there is the problem of the dynamically rearranging overlay. In particular, the main algorithm we later introduce in Section 4 will create a series of overlays $D_{1}, D_{2}, \ldots$ which will persist for only 2 rounds each. That means a node changes its position every 2 rounds. If now every node would keep all its routing messages and forward them from its new position, they would lose all the progress they made so far. Therefore, we define the so-called handover procedure using a helper graph $H_{i}$. For any point $p \in[0,1)$, let $S_{i}(p)$ be the swarm of $p$ in $D_{i}$ and $S_{i+1}(p)$ be the swarm of $p$ in $D_{i+1}$. Denote by $\left|S_{i}(p)\right|$ the number of nodes in the swarm $S_{i}(p)$. We assume that during the change from $D_{i}$ to $D_{i+1}$ each node from $S_{i}(p)$ can send a message to any set of nodes from $S_{i+1}(p)$, i.e., the nodes from a helper graph $H_{i}$ where the swarms $S_{i}(p)$ and $S_{i+1}(p)$ are adjacent. Formally, it is defined as follows:

Definition 5 (Handover Graph) Let $D_{0}, D_{1}, \ldots$ be a series of $L D S$ with $D_{i}=\left(V_{i}, E_{i}\right)$. Then the helper graph $H_{i}:=\left(V_{i} \cap V_{i+1}, E_{i}^{H}\right)$ is defined as follows:

$$
(v, w) \in E_{i}^{H} \Longleftrightarrow \exists p \in[0,1): v \in S_{i}(p) \wedge w \in S_{i+1}(p)
$$

Given this definition, we can easily see that the following holds:
Lemma 10 (Handover Property) Let $p \in[0,1)$ be an arbitrary point in the $[0,1]$-interval and let $S_{i+1}(p)$ be its swarm in $D_{i+1}$. Then in $H_{i}$ every node in $S_{i}^{r}(p)$ knows every other node in $S_{i+1}^{r}(p)$. The same holds for the left side.

Proof: As we will see, the property follows almost directly from the definition of swarms and the handover graph. We will prove the lemma only for the right side as the proof for the left side is completely analogous. Let $v$ be any node in $S_{i+1}^{R}(p)$. By the definition of the Handover graph every node in $S_{i}\left(p_{v}\right)$ knows $v$ as clearly $v \in S_{i+1}\left(p_{v}\right)$. Since $v$ is right of $p$ and within distance $\frac{c \lambda}{n}$ of $p$ as it is in $p^{\prime} s$ swarm, it holds that $S_{i}^{R}(p) \subset S_{i}(v)$. Thus, by combining it with the definition of the Handover Graph, we get that every node $S_{i}^{R}(p) \subset S_{i}(v)$ must know $v$. Since this holds for all nodes $v \in S_{i+1}^{R}(p)$, all nodes in $S_{i}^{R}(p)$ must know all nodes in $S_{i+1}^{R}(p)$ and the lemma follows.

Therefore, the switch can (almost) be handled like every other routing step from one swarm to another. The only difference is that we only send messages from and to the left and right side of each swarm respectively. However, if we choose the swarm size $\frac{c \lambda}{n}$ big enough, this makes no difference.

Later, in Section 4 we will see how to implement such a handover graph whereas here we just treat it as a property for a simpler description. Note that we call a helper graph $H_{i}$ good if for each $p \in[0,1)$ a $3 / 4$-fraction of all nodes in $S_{i+1}(p)$ is not churned out in the next round.

We summarize our observations in following definition for a routable series of graphs:
Definition 6 (Routable Graphs) Let $\mathcal{D}:=\left(D_{0}, H_{0}, D_{1}, H_{1} \ldots\right)$ be series of graphs defined on nodes $\mathcal{V}:=\left(V_{0}, V_{1}, \ldots\right)$, s.t., each $D_{i}$ consists of nodes in $V_{2 i}$. Then we call $\mathcal{D}$ routable, if

1. each $D_{i}$ is a $L D S$,
2. each $H_{i}$ enables is a handover from each $D_{i}$ to $D_{i+1}$, and
3. each $D_{i}$ (and $H_{i}$ ) is good, i.e., it holds $\left|S_{i}(p) \cap V_{2 i+1}\right| \geq 3 / 4 \cdot\left|S_{i}(p)\right|$ for all $p \in[0,1)$.

### 3.1 The Routing Algorithm

We now present a routing algorithm Alg-Routing for a dynamic series of routable graphs $\mathcal{D}:=$ $\left(D_{1}, H_{1}, D_{2}, H_{2}, \ldots\right)$. A trivial extension of the LDG routing algorithm would send each message to the whole swarm of each trajectory point. However, forwarding a message to a whole swarm would require $O\left(\log ^{2} n\right)$ messages to be sent in each step. In order to limit this to $O(\log n)$ messages ${ }^{7}$, we adapt the approach as follows. Assume a node $v \in V_{t}$ wants to route a message $m$. We first forward $m$ to all nodes in $S(v)$. Then, each node in $S(v)$ picks $r \in \Theta(1)$ nodes uniformly and independently at random from the next swarm $S\left(x_{1}\right)$ in the trajectory and forwards $m$ to them. Then, each node that received $m$ at least once, forwards it to $r$ nodes in $S\left(x_{2}\right)$ and so on. Only in the last step the message is forwarded to all nodes of the target swarm to ensure that the whole swarm receives the message. Listing 1 depicts the pseudo-code for Alg-Routing.

## Analysis

In this section, we analyze Alg-Routing. In particular, we prove that w.h.p. all messages reach their target and further analyze the dilation, i.e., the number of steps until a message reaches its target, and the congestion, i.e., the number of messages handled by each node in a round. Note that the latter depends on how many messages are sent each round and how their destinations are chosen. We would like to remark that we assume that each node sends exactly the same number of messages and chooses their destinations independently and uniformly at random.

Theorem 11 Let $\mathcal{D}$ be a routable series of $L D S$ defined on nodes $\mathcal{V}:=\left(V_{1}, V_{2}, \ldots\right)$. Further, let each node $v \in V_{1}$ start $t \in \mathbb{Z}_{+}$messages to random targets $p \in[0,1)$. Then Alg-Routing with $r \geq 16$ delivers each message with dilation exactly $2 \lambda+2$ and congestion $O(t \log n)$ w.h.p.

Proof: We begin the proof with the following lemma where we show that each message reaches its target swarm after exactly $2 \lambda+2$ rounds if it is not churned out.

Lemma 12 Let $\mathcal{D}$ be a routable series of $L D S$ defined on nodes $\mathcal{V}$. Let $v$ be any node in $V_{t} \in \mathcal{V}$, which sends a message to point $p \in[0,1)$ along the trajectory $\tau(v, p)$ using Alg-Routing. Then it holds that the message arrives at a node in $S_{\lambda+1+t}(p)$ in exactly $2 \lambda+2$ rounds.

Proof: The proof follows by an induction over the trajectory $i=0, \ldots, \lambda$ and the fact that by our choice of $\lambda$ the points $x_{\lambda}$ and $p$ are close by.

For the induction, observe that in each even step $j$ the message is forwarded along the trajectory and therefore moves from $S_{j}\left(x_{j-1}\right)$ to $S_{j}\left(x_{j}\right)$, whereas in each odd step the message is handed over and therefore moves from $S_{j}\left(x_{j}\right)$ to $\left.S_{j+1}\left(x_{j}\right)\right)$. Lemma 5 and the handover property imply that the nodes have necessary connections for each step but the last. We now prove by induction that, for each $i \in[\lambda]$, that each copy of the message is stored at a node $v \in S_{i}\left(x_{i}\right)$ in round $2 i+t$. W.l.o.g. assume the message is initiated in round $t=0$.
(IB) Consider round $i=0$, i.e., the round in which the message is started. In this round, the message is at $v=x_{0}$ and therefore, is known by $S_{0}\left(x_{0}\right)$ in the first step.
(IS) Now suppose that the induction hypothesis holds for any arbitrary $i \in[0, \lambda-1]$. We get that in round $2 i$ any copy of the message is at a node in $S_{i}\left(x_{i}\right)$. Now since the round $2 i$ is an even round, Alg-Routing performs a forwarding step in $D_{i}$ along the trajectory, i.e., every copy of the message is sent to some node in $S_{i}\left(x_{i+1}\right)$. Observe that, the swarm property in Lemma 5 ensures that each node in $S_{i}\left(x_{i}\right)$ has a connection to $S_{i}\left(\frac{x_{i}}{2}\right)$ and $S_{\lambda}\left(\frac{x_{i}+1}{2}\right)$. Now since $x_{i+1}$ is either $\frac{x_{\lambda}}{2}$ or $\frac{x_{\lambda}+1}{2}$, it holds that each node $S_{i}\left(x_{i}\right)$ has an edge to each node in $S_{i}\left(x_{i+1}\right)$. Therefore, every copy of the message can be successfully forwarded to each node in $S_{i}\left(x_{i+1}\right)$. Next, in round

[^6]$2 i+1$, Alg-Routing performs a handover operation on every copy of the message in $S_{i}\left(x_{i+1}\right)$ (overlay $H_{i}$ ). Now we use the Handover Property and observe that each node in $S_{i}\left(x_{i+1}\right)$ has by the definition of $H_{i}$, an edge to each node in $S_{i+1}\left(x_{i+1}\right)$. Therefore, every copy of the message can be successfully be sent to a node in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ and therefore, available in round $2 i+2:=2(i+1)$. This concludes the induction step.

The induction above implies that the message is known by all nodes in $S_{\lambda}\left(x_{\lambda}\right)$ in round $2 \lambda$. Now recall that $x_{\lambda}$ and $p$ are equal in their first $\lambda$ bits. This implies that the distance between $x_{\lambda}$ and $p$ is at most

$$
\begin{aligned}
d\left(x_{\lambda}, p\right) & \leq \frac{1}{2^{\lambda}} \\
& =\frac{1}{e^{\log (2) \lambda}}=\frac{1}{\left(e^{\lambda}\right)^{\log (2)}} \\
& =\frac{1}{(\kappa n)^{2 \log (2)}} \\
& \leq \frac{1}{\kappa n} \leq \frac{1}{n}
\end{aligned}
$$

as we defined $\lambda:=2 \log (\kappa n)$. Therefore, the swarms $S_{\lambda}\left(x_{\lambda}\right)$ and $S_{\lambda}(p)$ are adjacent and the message can be forwarded and handed over as described in the induction step. This proves the lemma.

In the proof of Lemma 12 we omitted the fact that not all nodes of a swarm forward the message as they may be churned out before they can do so. Of course, if a complete swarm is churned out, the message surely can't be forwarded along the trajectory. However, as stated earlier we assume that all swarms are good, i.e., only a constant fraction of each swarm is malicious and does not forward the message.

Lemma 13 Consider a set $S$ with $|S| \geq \frac{c \lambda}{2}$ nodes picked uniformly at random from the set of all nodes in any given round. Let $G \subset S$ be the set of good nodes in $S$, Then it holds that for churn parameters $\alpha=\frac{1}{16}$ and $\kappa=\left(1+\frac{1}{16}\right)$, with $c \geq 510 k$,

$$
\operatorname{Pr}\left[|G| \leq \frac{14}{17}|S|\right] \leq \frac{1}{n^{k}}
$$

Proof: Observe that by definition, for a churn rate of $\alpha$, there are at least $(1-\alpha) n$ nodes that would survive into the next round. Therefore, there are at least $\frac{15}{16} n$ good nodes in any given round. Also, since there could be at most $\kappa n$ nodes, there are at most $\frac{17}{16} n$ nodes in any given round. Therefore, the fraction of good nodes in the $[0,1)$-interval in any given round is then at least $\frac{15}{17}$. Furthermore, nodes pick their position uniformly at random in the $[0,1)$-interval.

Let $S \subset V_{t}$ be a set of nodes in round $t$ picked uniformly at random from the $[0,1)$-interval such that, $|S| \geq \frac{c \lambda}{2}$. For each $v \in S$, let $X_{v}$ be a $\{0,1\}$ r.v. such that $X_{v}=1$ if $v$ is a good node, and $X_{v}=0$ otherwise. We know that,

$$
\operatorname{Pr}\left[X_{v}=1\right]=\frac{15}{17}
$$

Therefore,

$$
\mathbf{E}\left[X_{v}\right]=\frac{15}{17}
$$

The expected number of good nodes in the set $S$ is then,

$$
\mathbf{E}\left[\sum_{v \in S} X_{v}\right]=\frac{15}{17}|S|
$$

The random variables $X_{1}, \ldots, X_{|S|}$ represent random sampling without replacement and therefore observe a permutation distribution which are negatively associated [12].

Now, let $X_{S}=\sum_{v \in S} X_{v}$ be the number of good nodes in the set $S$. Then applying the Chernoff bound on NA random variables with $c \geq 510 \cdot k$ we get,

$$
\operatorname{Pr}\left[X_{S} \leq\left(1-\frac{1}{15}\right) \mathbf{E}\left[X_{S}\right]\right] \leq \exp \left(-\frac{|S|}{255}\right) \leq \exp \left(-\frac{c \lambda}{510}\right) \leq \exp (-k \lambda)=n^{-k}
$$

Thus, as long as we observe only $O\left(n^{k}\right)$ swarms and each side of each swarm has more than $\frac{c \lambda}{2}$ nodes w.h.p, a simple union bound implies that in all swarms both the left and right side are good w.h.p.

Using Lemma 13 we can now show that the messages reach their destination w.h.p.
Lemma 14 Let $m$ be a message that is routed along $\tau(v, p):=x_{0}, \ldots, x_{\lambda+1}$ using Alg-Routing. Then, it holds that all nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ receive $m$ w.h.p after exactly $2 \lambda+2$ rounds for $r=16$.

Proof: We prove by induction that, for each $t \in[2 \lambda+2]$, it holds that if $t$ is an even round then, at least half of all good nodes in $S_{t / 2}\left(x_{t / 2}\right)$ receive the message $m$ w.h.p. Otherwise, if $t$ is an odd round, at least half of all good nodes in $S_{\lceil t / 2\rceil}^{R}\left(x_{\lceil t / 2\rceil-1}\right)$ and $S_{\lceil t / 2\rceil}^{L}\left(x_{\lceil t / 2\rceil-1}\right)$ receive the message $m$ w.h.p.
(IB) Consider round $t=0$, i.e., the round in which the message is initiated. If this is an even round, then observe that Alg-Routing forwards message $m$ from node $v$, i.e., $x_{0}$, to all nodes in $S_{0}\left(x_{0}\right)$. Using Lemma 6 and 13 we can conclude that at least half of all good nodes in the swarm $S_{0}\left(x_{0}\right)$ received the message $m$ and survive until the next round. Therefore, the induction hypothesis holds. Otherwise, if $t=0$ is an odd round, Alg-Routing handovers the message $m$ from node $v$ to all nodes in $S_{1}\left(x_{0}\right)$. Again, using Lemma 6 and 13, $S_{1}\left(x_{0}\right)$ is a good swarm and therefore, the induction hypothesis holds.
(IS) Now suppose the induction hypothesis holds for any arbitrary $t \in[2 \lambda+1]$ and w.l.o.g. assume that round $t=2 \lambda+1$ is an odd round.
Note that the handover step works analogously, the only difference is that we do not send messages from one swarm to another but from the right/left side from one swarm to the right/left side of another. However, since these sets are smaller, we assume the swarm size to be at least $\frac{c \lambda}{n}$ (and not $\left.\frac{2 c \lambda}{n}\right)$. This way, we cover both cases. By the induction's hypotheses, at least half of all good nodes in $S_{\lambda+1}\left(x_{\lambda}\right)$ received message $m$ w.h.p. and therefore, each of these nodes forward $r$ copies of the message $m$ to nodes picked uniformly and independently at random from the swarm $S_{\lambda+1}\left(x_{\lambda+1}\right)$. We now show that at least half of all good nodes in the swarm $S_{\lambda+1}\left(x_{\lambda+1}\right)$ receive the message w.h.p. in the round $2 \lambda+2$.
From Lemma 6] it holds that $\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right| \leq 4\left|S_{\lambda+1}\left(x_{\lambda}\right)\right|$ w.h.p. Furthermore, by the induction hypothesis we know that at least half of all good nodes in $S_{\lambda+1}\left(x_{\lambda}\right)$ received $m$ and therefore forward $r$ copies of $m$. Thus, in total there are at least $r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|$ copies of $m$ sent to $S_{\lambda+1}\left(x_{\lambda+1}\right)$. The probability that any of these messages is sent to a given node $v^{\prime} \in S_{\lambda+1}\left(x_{\lambda+1}\right)$ is $\frac{1}{\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|}$, since the destinations are chosen uniformly at random. Observe that one can view the forwarding of messages from $S_{\lambda+1}\left(x_{\lambda}\right)$ to uniformly and independently picked nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ as a balls-into-bins experiment, where (at least) $r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|$ balls are thrown into $\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|$ bins. Using Propositions 2 and 3 one can show that the number of nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ that receive at least one ball is NA [7]. For $i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$ and $j \in$ $\left\{1, \ldots, r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$, let $X_{i, j}$ be an indicator random variable such that, $X_{i, j}=1$ if message $j$ is sent to node $i$ (picked uniformly and independently at random) by Alg-Routing in round $2 \lambda+1$, and $X_{i, j}=0$ otherwise.

Lemma 15 (Zero-One Lemma [7]) If $Y_{1}, \ldots, Y_{n}$ are zero-one random variables such that $\sum_{i} Y_{i}=$ 1 , then $Y_{1}, \ldots, Y_{n}$ are $N A$.

For any fixed $j \in\left\{1, \ldots, r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$, let $Y_{i}:=X_{i, j}$ for all $i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right.$. Then, from Lemma 15 we know that the random variables $Y_{1}, \ldots, Y_{\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|}$ are NA. Since each message $j \in\left\{1, \ldots, r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$ is destined to a node that is picked uniformly and independently at random, using Proposition 2, we can conclude that the set of random variables $\left(X_{i, j}\right)_{i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}, j \in\left\{1, \ldots, r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}}$ are NA.
For each $i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$ consider a non-decreasing function as follows,

$$
X_{i}= \begin{cases}1 & \sum_{j \in\left\{1, \ldots, r / 8 \cdot\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right.} X_{i j}>0 \\ 0 & \text { otherwise }\end{cases}
$$

Therefore, for each $i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}$,

$$
\begin{aligned}
\operatorname{Pr}\left[X_{i}=1\right] & =1-\left(1-\frac{1}{\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|}\right)^{\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right| \cdot r / 8} \\
& \geq 1-\left(\frac{1}{e}\right)^{\frac{r}{8}}
\end{aligned}
$$

From Proposition 3, we know that the random variables $X_{1}, \ldots, X_{\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|}$ are NA.

Let $G_{\lambda+1}\left(x_{\lambda+1}\right) \subseteq S_{\lambda+1}\left(x_{\lambda+1}\right)$ denote the set of good nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$. For each $v \in$ $G_{\lambda+1}\left(x_{\lambda+1}\right), G_{v}$ be a $\{0,1\}$ random variable such that, $G_{v}=1$ if node $v$ received at least one copy of $m$ from some node in $S_{\lambda+1}\left(x_{\lambda}\right)$ and $G_{v}=0$, otherwise. Observe that $X_{i}$ denotes if a node in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ received at least one message in round $2 \lambda+2$ when exactly $\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right| \cdot r / 8$ messages are sent to $S_{\lambda+1}\left(x_{\lambda+1}\right)$ uniformly and independently at random. Since $\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right| \cdot r / 8$ is a lower bound on the number of message being sent, for each $v \in G_{\lambda+1}\left(x_{\lambda}+1\right)$ we have that,

$$
\operatorname{Pr}\left[G_{v}=1\right] \geq \operatorname{Pr}\left[X_{v}=1\right]
$$

Moreover, the random variables $\left(G_{v}\right)_{v \in G_{\lambda+1}\left(x_{\lambda+1}\right)}$ are also NA as they can be seen as a subset of $\left(X_{v}\right)_{v \in S_{\lambda+1}\left(x_{\lambda+1}\right)}$. To see this, first recall that the random variables $\left(X_{i}\right)_{i \in\left\{1, \ldots,\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right\}}$ are NA. In particular, this fact is independent of the number of messages and nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$. Now observe that for each $v \in G_{\lambda+1}\left(x_{\lambda+1}\right)$, the random variable $G_{v}$ is a non-decreasing function of its associated random variable $X_{v}$. Therefore, we can conclude that the random variables $\left(G_{v}\right)_{v \in G_{\lambda+1}\left(x_{\lambda+1}\right)}$ are also NA.
Let $G:=\sum_{v \in G_{\lambda+1}\left(x_{\lambda+1}\right)} G_{v}$ be a random variable that counts the number of good nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ that received at least one copy of $m$ in round $2 \lambda+2$. From Lemma 13, we know that there are $14 / 17$ fraction of good nodes in $S_{\lambda+1}\left(x_{\lambda+1}\right)$ w.h.p. Therefore, the expected number of good nodes that receive at least one message is given by,

$$
\begin{aligned}
\mathbf{E}[G] & =\sum_{v \in G_{\lambda+1}\left(x_{\lambda+1}\right)} \mathbf{E}\left[G_{v}\right] \\
& \geq \sum_{v \in G_{\lambda+1}\left(x_{\lambda+1}\right)}\left(1-\frac{1}{e^{r / 8}}\right) \\
& =\left(1-\frac{1}{e^{r / 8}}\right) \frac{14}{17}\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right| .
\end{aligned}
$$

To complete the induction it suffices to show that,

$$
\operatorname{Pr}\left[G \leq \frac{1}{2}\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right] \leq \frac{1}{n^{k}}
$$

Using $\delta=\frac{2}{7}$ in Lemma 4 with $r=16$ we get,

$$
\operatorname{Pr}\left[\left.G \leq\left(1-\frac{2}{7}\right)\left(1-\frac{1}{e^{2}}\right) \frac{14}{17} \right\rvert\, S_{\lambda+1}\left(x_{\lambda+1}\right)\right] \leq \exp \left(-\frac{4}{49 \cdot 3}\left(1-\frac{1}{e^{2}}\right) \frac{14}{17}\left|S_{\lambda+1}\left(x_{\lambda+1}\right)\right|\right)
$$

Lemma 6 then implies,

$$
\begin{aligned}
& \leq \exp \left(-\frac{4}{49 \cdot 3}\left(1-\frac{1}{e^{2}}\right) \frac{14}{17} \frac{c \lambda}{2}\right) \\
& =\exp \left(-\frac{4}{357}\left(1-\frac{1}{e^{2}}\right) c \lambda\right)
\end{aligned}
$$

for $c \geq 510 \cdot k$ we get,

$$
\leq n^{-k}
$$

We conclude the analysis by observing each node's congestion. Therefore, we first bound the expected number of trajectories that cross an interval in each round. Note that a trajectory is defined on points in $[0,1)$ and not on actual nodes (except the first and last element). We can see that it holds:

Lemma 16 Assume all nodes choose their position independently and uniformly at random in the $[0,1)$ interval. Moreover, let each node send $\varphi \in \mathbb{Z}_{+}$messages to targets picked independently and uniformly at random from the $[0,1)$-interval. Then for every $I \subset[0,1)$ it holds that,

1. $X_{I}^{j}$ is the sum of independent $\{0,1\}$ random variable,
2. $\mathbf{E}\left[X_{I}^{j}\right]=k n|I|$,
where $X_{I}^{j}$ is a random variable that counts the number of trajectories that have their $j^{\text {th }}$ step in the interval $I$ and $|I|$ denotes the size of the interval $I$.

Proof: W.l.o.g. assume that $I:=[a, b]$ with $0 \leq a \leq b<1$.

1. Let $X_{I:(v, i)}^{j}$ be a $\{0,1\}$ random variable such that, $X_{I:(v, i)}^{j}=1$ if the trajectory of a message $i \in[\varphi]$ started by a node $v$ crosses the interval $I$ in its $j^{t h}$ step and $X_{I:(v, i)}^{j}=0$, otherwise. Then the number of messages with their $j^{t h}$ step in the interval $I$ is given by,

$$
X_{I}^{j}=\sum_{v \in V} \sum_{i \leq \varphi} X_{I:(v, i)}^{j}
$$

Observe that a message's trajectory is uniquely defined by the starting node $v$ and the end point $p$. Since for each message the target node is chosen uniformly and independently at random from $[0,1)$-interval, we conclude that the set of random variables $\left(X_{I:(v, i)}^{j}\right)_{v \in V, i \in[\varphi]}$ are independent.
2. We prove by induction that for each $j \in[\lambda+1]$ step along the trajectory, it holds that,

$$
\mathbf{E}\left[X_{I}^{j}\right]=k n|I|
$$

Consider step $j=0$ i.e., the step where is messages are at their starting node at $x_{0}$. For each $v \in V$, let $X_{v}$ be a $\{0,1\}$ random variable such that, $X_{v}=1$ if node $v \in I$ and $X_{v}=0$, otherwise. Observe that the nodes pick their positions uniformly and independently at random in the $[0,1)$-interval. Therefore,

$$
\operatorname{Pr}\left[X_{v}=1\right]=|I| .
$$

This implies,

$$
\mathbf{E}\left[\sum_{v \in V} X_{v}\right]=\sum_{v \in V} \operatorname{Pr}\left[X_{v}=1\right]=n|I| .
$$

Since each node initiates $\varphi \in \mathbb{Z}_{+}$messages to randomly picked targets in the [0,1)-interval,

$$
\mathbf{E}\left[X_{I}^{0}\right]=\mathbf{E}\left[\sum_{v \in V} \varphi \cdot X_{v}\right]=n \varphi|I|
$$

Now assume the induction hypothesis holds for every $j \in[\lambda]$, this implies,

$$
\mathbf{E}\left[X_{I}^{\lambda}\right]=\varphi n|I|
$$

Let $I_{0}:=I \cap[0,1 / 2)$ and $I_{1}:=I \cap[1 / 2,1)$ be the parts of $I$ that lie in the first and second half of $[0,1)$-interval, respectively. Then,

$$
\mathbf{E}\left[X_{I}^{\lambda+1}\right]=\mathbf{E}\left[X_{I_{0}}^{\lambda+1}\right]+\mathbf{E}\left[X_{I_{1}}^{\lambda+1}\right]
$$

Observe that for each $i \in\{0,1\}$ the bit representation of each point on $I_{i}$ begins with $i$. Therefore for any message $m$ destined to a uniformly and independently picked target $p \in[0,1)$-interval, it's trajectory crosses interval $I_{i}$ in the $j^{\text {th }}$ step of the trajectory if and only if $p_{j}=i$, where $p_{j}$ is the $j^{\text {th }}$ most significant bit of $p$ in the binary representation. W.l.o.g. we only analyze $I_{0}$ to show that,

$$
\mathbf{E}\left[X_{I_{0}}^{\lambda+1}\right]=\varphi n\left|I_{0}\right|
$$

The proof for $\mathbf{E}\left[X_{I_{1}}^{\lambda+1}\right]=\varphi n\left|I_{1}\right|$ is analogous and follows using similar arguments.
Consider an arbitrary trajectory $\tau:=x_{0}, \ldots, x_{\lambda+1} \in[0,1)^{\lambda+2}$ with $x_{\lambda+1} \in I_{0}:=\left[a, b_{0}\right]$. From Definition 4 it must then hold that $x_{\lambda}=2 x_{\lambda+1}$. Therefore, for all trajectories that cross the interval $I_{0} \in[0,1 / 2)$ it holds that $x_{\lambda}=2 x_{\lambda+1}$ and hence, the $\lambda^{\text {th }}$ step must be in the interval $J:=\left[2 a, 2 b_{0}\right]$. The size of the interval is then,

$$
|J|=2\left|I_{0}\right|
$$

By the induction hypothesis we know that,

$$
\mathbf{E}\left[X_{J}^{\lambda}\right]=\varphi n|J|=2 k n\left|I_{0}\right|
$$

Note that since for every message it's target is picked uniformly and independently at random from the $[0,1)$-interval, this is equivalent to the thought experiment of flipping a fair coin for each bit of the target address. Therefore, the probability that a trajectory $\tau$ in interval $J$ points towards the interval $I_{0}$ in it's $(\lambda+1)^{\text {th }}$ step is then,

$$
\operatorname{Pr}\left[\tau_{\lambda+1} \in I_{0}\right]=\frac{1}{2}
$$

where $\tau_{\lambda+1}$ denotes the position of the trajectory $\tau$ in step $\lambda+1$.
Then, the expected number of trajectories that point from the interval $J$ to the interval $I_{0}$ is given by,

$$
\begin{aligned}
\mathbf{E}\left[X_{I_{0}}^{\lambda+1}\right] & =\sum_{\ell=0}^{\infty} \mathbf{E}\left[X_{I_{0}}^{\lambda+1} \mid X_{J}^{\lambda}=\ell\right] \cdot \operatorname{Pr}\left[X_{J}^{\lambda}=\ell\right] \quad \text { (Law of Total Expectation) } \\
& =\sum_{\ell=0}^{\infty} \sum_{\tau \in[1, \ldots, \ell]} \operatorname{Pr}\left[\tau_{\lambda+1} \in I_{0}\right] \cdot \operatorname{Pr}\left[X_{J}^{\lambda}=\ell\right] \\
& =\sum_{\ell=0}^{\infty} \frac{\ell}{2} \cdot \operatorname{Pr}\left[X_{J}^{\lambda}=\ell\right]
\end{aligned}
$$

$$
\begin{aligned}
& =\frac{1}{2} \sum_{\ell=0}^{\infty} \ell \cdot \operatorname{Pr}\left[X_{J}^{\lambda}=\ell\right] \\
& =\frac{1}{2} \mathbf{E}\left[X_{J}^{\lambda}\right]=k n\left|I_{0}\right| .
\end{aligned}
$$

Since $|I|=\left|I_{0}\right|+\left|I_{1}\right|$ we get,

$$
\mathbf{E}\left[X_{I}^{\lambda+1}\right]=\varphi n\left|I_{0}\right|+\varphi n\left|I_{1}\right|=\varphi n|I| .
$$

This completes the induction.
Using Lemma 16 we can then bound from above the expected congestion for Alg-Routing.
Lemma 17 If each node picks its position and the $t \in \mathbb{Z}_{+}$target nodes uniformly and independently at random, then Alg-Routing has a expected congestion at most $24 r t c \lambda$.

Proof: Let $v \in V$ be any node and let $\left[v \pm \frac{c \lambda}{n}\right]=: I_{v} \subset[0,1)$ be an interval that contains all points $p$ with $v \in S(p)$. Observe that a message may be routed via $v$ only if its trajectory passes the interval $I_{v}$. From Lemma 16 we know that for any given round $j$ the expected number of trajectories that cross the interval $I_{v}$ is then,

$$
\mathbf{E}\left[X_{I_{v}}^{j}\right]=2 t c \lambda
$$

where $X_{I_{v}}^{j}$ is a random variable that counts the number of trajectories that have their $j^{\text {th }}$ step in the interval $I_{v}$. Using the Chernoff bound we get,

$$
\begin{aligned}
\operatorname{Pr}\left[X_{I_{v}}^{j} \geq 3 t c \lambda\right] & \leq \exp \left(-\frac{1}{6} t c \lambda\right) \\
\text { for } c \geq 510 \cdot k, & \\
& \leq n^{-k}
\end{aligned}
$$

Therefore, the total number of trajectories that pass interval $I_{v}$ in any round $j$ is at most $3 t c \lambda$ w.h.p. We know from Definition 4 that a trajectory passing interval $I_{v}$ in step $j$, had it's step $j-1$ in either interval $J^{0}:=\left[2\left(v \pm \frac{c \lambda}{n}\right)\right]$ with $p_{j}=0$, or in the interval $J^{1}:=\left[2\left(v \pm \frac{c \lambda}{n}\right)-1\right]$ with $p_{j}=1$. Observe that the size of these intervals i.e., $\left|J^{0}\right|=\left|J^{1}\right|=2\left|I_{v}\right|$. From Lemma 6 we know that the number of nodes in each of these interval i.e., $J^{0}$ and $J^{1}$, are at most $8 c \lambda$ w.h.p. Therefore, the total number of messages that will be forwarded to the interval $I_{v}$ in any given round is then at most $M=(r \cdot 8 c \lambda) \cdot(3 t c \lambda)$ w.h.p.

Let $\left(X_{i}^{v}\right)_{i \in[M]}$ be a set of $\{0,1\}$ random variables such that, $X_{i}^{v}=1$ if message $m_{i}$ is sent to node $v$ and $X_{i}^{v}=0$, otherwise. From Lemma 6 we know that any interval of size $\frac{2 c \lambda}{n}$ has at least $c \lambda$ nodes w.h.p. Then,

$$
\operatorname{Pr}\left[X_{i}^{v}=1\right] \leq \frac{1}{c \lambda}
$$

The expected congestion is then,

$$
\mathbf{E}\left[\sum_{i \in[M]} X_{i}^{v}\right] \leq \frac{1}{c \lambda}(r \cdot 8 c \lambda) \cdot(3 t c \lambda)=24 r t c \lambda .
$$

Theorem 11 follows directly from Lemmas 1214 and 17 .

```
1 Desc: This algorithm is executed on a routable graph \(\mathcal{D}:=\left(D_{1}, H_{1}, \ldots\right)\). It routes a
            message \(m\) from any node \(u \in V_{t}\) to a node \(v\) (almost) uniformly picked from
        \(V_{t+2 \lambda+2}\).
2
3 Send a message m to a random node \(v \in V_{t+2 \lambda+2}\)
\(4 \quad p \longleftarrow\) Uniformly chosen from \([0,1)\)
    \(\Delta \longleftarrow\) Uniformly chosen from \([0,2 c \lambda]\)
    Route message \((m, p, \Delta)\) to target \(S(p)\) using Alg-Routing
7
8 Upon receiving ( \(m, p, \Delta\) ) from Alg-Routing
    \(P \longleftarrow\left\{w \in S(p) \mid p_{w} \in S(p)\right\}\)
    Choose \(w\) such that \(|\{u \mid u \in\langle p, w\rangle\}|=\Delta \bmod P\)
    Deliver \(m\) to \(w\)
```


### 3.2 The Random Sampling Algorithm

Besides routing to a random swarm $S(p)$ for some $p \in[0,1)$ the algorithm Alg-Routing can also be extended to send a message to a node (and not swarm) chosen uniformly at random. We call this algorithm Alg-Sampling. The underlying approach is adapted from King and Saia 14 and King et al. [13]. In their algorithm, King and Saia condition on the fact that the hash function that provides the nodes with their positions has certain properties. In particular, these properties are fulfilled with probability at least $1-\frac{3}{n}$ on a randomly chosen hash function $h$. This is not good enough for our case as we need a sampling routine that works correctly w.h.p., i.e., with probability $1-\frac{1}{n^{k}}$ where we can freely choose some $k>1$. We strongly believe that their approach can be adapted to work, w.h.p., however we will only show a weaker statement that suffices for our needs. In particular, we only need the sampling probabilities to be within a constant factor, i.e., between, $\frac{1}{4 n}$ and $\frac{5}{n}$. Our approach works as as follows. A node first picks a value $p \in[0,1)$ uniformly at random and routes the message to the swarm $S(p)$ using algorithm Alg-Routing. Then, the message is only delivered to some randomly chosen node $w \in S(p)$. For this, the message includes a random chosen number $\Delta \in[0,2 c \lambda]$. The message is then delivered to the node for which it holds $|\{u \mid u \in\langle p, w\rangle\}|=\Delta \bmod |S(p)|$, i.e., the $\Delta^{t h}$ node (in $S(p)$ ) that is right of $p$. Since all nodes in $S(p)$ know $|S(p)|$ (because they know the IDs of all nodes in $S(p)$ ) and $\Delta$ (because they all received the message) this can checked locally without further messages. The following lemma bounds the sampling probability.
Lemma 18 Let $\mathcal{D}$ be routable. Assume, a node $v \in V_{t}$ starts a message $m_{v}$ using Alg-Sampling. Then for all $u \in V_{t+2 \lambda+2}$,

$$
\operatorname{Pr}\left[u \text { receives } m_{v}\right] \in\left[\frac{1}{4 n}, \frac{5}{n}\right] .
$$

Furthermore, for any two nodes $v, w \in V_{t}$ that start messages $m_{v}$ and $m_{w}$, respectively, it holds that for all $u \in V_{t+2 \lambda+2}$,

$$
\operatorname{Pr}\left[u \text { receives } m_{v}\right]=\operatorname{Pr}\left[u \text { receives } m_{w}\right] .
$$

Proof: Let $Y(v, u)$ be a $\{0,1\}$ random variable such that $Y(v, u)=1$ in the event that $v$ samples $u$ i.e., the message $m_{v}$ is delivered to $u$. Further, let $p \in[0,1)$ be the point that $v$ chooses in line 4 of Alg-Sampling. Then, we make the following observations.

1. As a necessary condition to sample the node $u$, node $v$ must pick some point $p \in S(u)$. Otherwise, $u$ will never be considered in the second step. This happens with probability $\frac{2 c \lambda}{n}$.
2. Given that $p \in S(u)$, Alg-Sampling still needs to pick $u$ uniformly from all nodes in $S(p)$. The probability for this depends only on $|S(p)|$ and $\Delta$. By Lemma we know that w.h.p. it contains
at most $2 c \lambda$ nodes and at least $c \lambda / 2$ nodes. Since $\Delta \leq 2 c \lambda$ and $\frac{c}{2} \lambda \leq\left|S_{l}(p)\right| \leq 2 c \lambda$ there are at least one and at most 4 choices of $\Delta$ that result in $u$ being picked.

These observations are sufficient to prove the statement. For the lower bound we get that:

$$
\begin{aligned}
\operatorname{Pr}[Y(v, u)] & =\operatorname{Pr}\left[p \in S^{L}(u)\right] \cdot \operatorname{Pr}\left[\mathrm{u} \text { is picked from } S^{R}(p)\right] \\
& \geq \operatorname{Pr}\left[p \in S^{L}(u)\right] \cdot\left(\operatorname{Pr}\left[\left|S^{R}(p)\right| \leq 2 c \lambda\right] \frac{1}{2 c \lambda}+\operatorname{Pr}\left[\left|S^{R}(p)\right|>2 c \lambda\right] \frac{1}{n}\right) \\
& =\frac{c \lambda}{n} \cdot\left(\left(1-\frac{1}{n^{k}}\right) \frac{1}{2 c \lambda}+\frac{1}{n^{k+1}}\right) \\
& \geq \frac{1}{2 n}-\frac{1}{2 n^{k+1}} \\
& \geq \frac{1}{4 n} .
\end{aligned}
$$

The proof for the upper bound is analogous, we simply need to replace the upper bound for the swarm size with the lower bound. We get

$$
\begin{aligned}
\operatorname{Pr}[Y(v, u)] & =\operatorname{Pr}\left[p \in S^{L}(u)\right] \cdot \operatorname{Pr}\left[\mathrm{u} \text { is picked from } S^{R}(p)\right] \\
& \leq \operatorname{Pr}\left[p \in S^{L}(u)\right] \cdot\left(\operatorname{Pr}\left[\left|S^{R}(p)\right| \geq \frac{c}{2} \lambda\right] \frac{1}{c \lambda}+\operatorname{Pr}\left[\left|S^{R}(p)\right| \leq \frac{c}{2} \lambda\right] \cdot 1\right) \\
& =\frac{2 c \lambda}{n} \cdot\left(\left(1-\frac{1}{n^{k}}\right) \frac{1}{c \lambda}+\frac{1}{n^{k}}\right) \\
& \leq \frac{c \lambda}{n} \cdot\left(\frac{4}{c \lambda}+\frac{1}{n^{c}}\right) \\
& =\frac{4}{n}+\frac{2 c \lambda}{n^{k+1}} \\
& \leq \frac{5}{n} .
\end{aligned}
$$

This concludes the first part of the lemma.
It remains to show that (messages of) nodes $v$ and $w$ are delivered to a node $u$ with the same probability. Let $p_{v}$ and $p_{w} \in[0,1)$ be the points picked by these nodes for their respective messages. Further, let $\Delta_{v}$ and $\Delta_{w} \in[0,2 c \lambda]$ be random numbers. Let $\ell(u, p):=|\{w \mid w \in\langle p, u\rangle\}|$ and $\mathcal{U}(u, p)$ be the number of possible choices of $\Delta$ that lead to $u$ being picked given that the message is routed to $p$ in the first step. Note that any point $p \in[0,1)$ and any $\Delta \in[0,2 c \lambda]$ is picked with equal probability by $v$ and $w$. Therefore, it holds that,

$$
\begin{aligned}
\operatorname{Pr}[Y(v, u)] & =\operatorname{Pr}\left[p_{v} \in S^{L}(u)\right] \cdot \operatorname{Pr}\left[\mathrm{u} \text { is picked from } S^{R}\left(p_{v}\right)\right] \\
& =\operatorname{Pr}\left[p_{v} \in S^{L}(u)\right] \cdot\left(\sum_{p^{*} \in S^{L}(u)} \operatorname{Pr}\left[p_{v}=p^{*}\right] \cdot \operatorname{Pr}\left[\ell\left(u, p^{*}\right)=\Delta_{v} \bmod \left|S^{R}\left(p^{*}\right)\right|\right]\right) \\
& =\operatorname{Pr}\left[p_{v} \in S^{L}(u)\right] \cdot\left(\sum_{p^{*} \in S^{L}(u)} \operatorname{Pr}\left[p_{v}=p^{*}\right] \frac{\mathcal{U}\left(p^{*}\right)}{\left|S^{R}\left(p^{*}\right)\right|}\right) \\
& =\operatorname{Pr}\left[p_{w} \in S^{L}(u)\right] \cdot\left(\sum_{p^{*} \in S^{L}(u)} \operatorname{Pr}\left[p_{w}=p^{*}\right] \frac{\mathcal{U}\left(p^{*}\right)}{\left|S^{R}\left(p^{*}\right)\right|}\right) \\
& =\operatorname{Pr}\left[p_{w} \in S^{L}(u)\right] \cdot\left(\sum_{p^{*} \in S^{L}(u)} \operatorname{Pr}\left[p_{w}=p^{*}\right] \cdot \operatorname{Pr}\left[\ell\left(u, p^{*}\right)=\Delta_{w} \bmod \left|S^{R}\left(p^{*}\right)\right|\right]\right)
\end{aligned}
$$

$$
\begin{aligned}
& =\operatorname{Pr}\left[p_{w} \in S^{L}(u)\right] \cdot \operatorname{Pr}\left[\mathrm{u} \text { is picked from } S^{R}\left(p_{w}\right)\right] \\
& =\operatorname{Pr}[Y(w, u)] .
\end{aligned}
$$

The third equality is due to the fact that,

$$
\operatorname{Pr}\left[p_{v} \in S^{L}(u)\right]=\operatorname{Pr}\left[p_{w} \in S^{L}(u)\right]=\frac{c \lambda}{n}
$$

and for any $p \in[0,1)$,

$$
\operatorname{Pr}\left[p_{v}=p\right]=\operatorname{Pr}\left[p_{w}=p\right] .
$$

This follows from the fact that all nodes use the same random hash function $h$ to compute the positions. This concludes the proof.

## 4 The Maintenance Algorithm

In this section we present our main contribution. Our algorithm maintains a routable dynamic overlay $\mathcal{D}=\left(D_{0}, H_{0}, D_{1}, \ldots\right)$ with high probability. Before we present the algorithm, we will first give an overview on our assumptions and choice of churn parameters.

We assume a $(2,2 \lambda+7)$-late adversary with a churn rate $(n / 16,2 \lambda+7)$. This implies that within $2 \lambda+7$ (i.e. $O(\log n)$ ) rounds, a constant fraction of nodes can be subjected to churn. Further, we assume that the number of nodes in any round is at most $(1+1 / 16) n$. Note that the values $\alpha=1 / 16$ and $\kappa=(1+1 / 16)$ are chosen for the sake of convenience in analysis. We require a bootstrap phase of length $B:=2 \lambda+2$ at the beginning of the algorithm. In this phase no churn occurs and this enables us to initialize our algorithm. Such a bootstrap phase is a standard assumption in the area of churn-resistant overlays and is very likely necessary to construct a robust overlay.

We also assume that the system starts in an initial LDS $D_{0}$ in round 0 . This assumption is made for convenience as the initial overlay can easily be constructed in the churn-free bootstrap phase using algorithms from [10, 11]. Using their techniques this can be achieved in $O\left(\log ^{2} n\right)$ rounds with a deterministic algorithm or in $O(\log n)$ rounds w.h.p. with a randomized algorithm. Both these algorithms assume that the congestion and degree of each node is polylogarithmic, so they fit well into our computational model. We would like to remark that since our focus lies on fast reconfiguration and not on optimizing the bootstrap phase we omit the algorithmic details and the corresponding analysis. For ease of notation we will refer to round $t+B$ simply as $t$.

Let $\overline{V_{t}}:=V_{t} \cap V_{t-1}$ denote the set of all rounds except for the newly joined nodes in any given round $t$. Over the course of this section we distinguish between three types of nodes in each round $t$. Namely, the set of mature nodes $M_{t} \subseteq \bar{V}_{t}$, which are nodes that are in the network for at least $2 \lambda+2$ rounds (or $2 \lambda+3$ if they joined in an odd round), the set of fresh nodes $F_{t}:=\bar{V}_{t} \backslash M_{t}$ which are nodes that are at least one round, but less than $2 \lambda+2$ rounds old, and the set of newly joined nodes i.e., $V_{t} \backslash \bar{V}_{t}$. Observe that $\bar{V}_{t}:=M_{t} \cup F_{t}$ and due to our choice of churn parameters, it holds $M_{t} \geq n\left(1-\frac{1}{16}\right)$ and $F_{t} \leq n / 16$.

Our algorithm is comprised of two subroutines, Alg-LDS and Alg-Random, that are executed in concurrently. Alg-LDS ensures that all mature nodes build a routable overlay each round and ALg-Random makes sure that all fresh nodes (which are not part of the routable overlay) stay connected to $O(\log n)$ mature nodes until they mature themselves. This ensures that the matures node can route a message on behalf of the fresh nodes over the overlay. The main result of this section (and this paper) is stated in the following theorem.

Theorem 19 Subroutines Alg-LDS and Alg-Random maintain a series of overlays $\mathcal{D}$ such that for $O\left(n^{k}\right)$ rounds w.h.p,

1. the mature nodes form a routable series of graphs $\mathcal{D}:=\left(D_{0}, H_{0}, D_{1}, \ldots\right)$,
2. each fresh node is known by $\Theta(\log n)$ mature nodes, and
3. the congestion is $O\left(\log ^{3} n\right)$ per node and round.

We would like to remark that both Alg-LDS and Alg-Random are heavily randomized and can possibly fail to create a connected and routable overlay if they are executed for too long. For example, the algorithm could fail if certain swarms are too small and/or too many messages are dropped by the routing algorithm. In these cases, the algorithm cannot construct the desired overlays and needs to be restarted along with another bootstrap phase. Given that the algorithm runs correctly w.h.p., i.e. no failure happens with probability $O\left(n^{-k}\right)$ for a tunable constant $k>0$, we can only guarantee that the algorithm runs smoothly for $O\left(n^{k}\right)$ rounds w.h.p. until some failure happens. This follows by a simple application of the union bound. Throughout this chapter we assume that $\frac{k \log (n)}{n} \ll 1$, i.e., both $k$ and $\log (n)$ are very small compared to $n$. Therefore, our algorithms only become applicable for large values of $n$ (say $n>10^{6} k$ ). However, as our goal is to show that the messages per node stays logarithmic in $n$ even under heavy churn, we believe that it is justified to only consider very high values of $n$. In particular, we do not seek (nor claim) that we make the optimal choice of constants.

Alg-LDS and Alg-Random exchange four types of messages between the nodes to build a series of overlay.

1. The message $\operatorname{Connect}(v)$ is sent by a fresh node $v$ to advertise itself to a mature node in the overlay. It only contains $v$ 's identifier.
2. The message $\operatorname{Create}\left(v, p_{v}^{t}\right)$ is used to introduce a node $v$ to its neighbours in $S_{t}\left(p_{v}^{t}\right)$. The message contains the node's ID and its position $p_{v}^{t}$ in the overlay $D_{t} \in[0,1)$.
3. The message $\operatorname{Join}\left(v, p_{v}^{t}\right)$ is used to introduce a node $v$ to nodes in $S_{t-1}\left(p_{v}^{t}\right)$. It is routed from its origin to position $p_{v}^{t}$ in overlay $D_{t}$. It contains the identifier of $v$ and its position $p_{v}^{t} \in[0,1)$.
4. The message $\operatorname{Token}(v)$ is sent by a mature node $v$ to a point in the $[0,1$ )-interval picked (almost) uniformly at random. It only contains $v$ 's identifier.

## Building a Routable Overlay

After the bootstrap phase, the algorithm ALG-LDS creates a series of overlays $\mathcal{D}=\left(D_{0}, H_{0}, D_{1}, H_{1}, \ldots\right)$ that contain all mature nodes in any given round. In particular, in every even round $t=2 i$ the algorithm creates an LDS $D_{i}$ which consists of all mature nodes $\bar{V}_{t}$. In each odd round $t=2 i+1$ the algorithm creates a handover graph $H_{i}$ in which for each $p \in[0,1)$ it holds that $S_{i}(p)$ and $S_{i+1}(p)$ are adjacent (Lemma 5), where $S_{i}(p)$ and $S_{i+1}(p)$ are the swarms of point $p$ in $D_{i}$ and $D_{i+1}$, respectively.

To construct a series of overlays, all mature nodes continuously choose new positions $p_{v}^{0}, p_{v}^{1}, \ldots$ for the corresponding overlays $D_{0}, D_{1}, \ldots$ and use the routing algorithm to find their neighbors. More precisely, the construction of $D_{i}$ begins in round $2 i-(2 \lambda+2)$. Every node $v \in \bar{V}_{2(i-\lambda-1)}$ picks a position $p_{v}^{i} \in[0,1]$ uniformly at random and routes its $I D$ to the target $p_{v}^{i}$ along the trajectory $\left(v, x_{1}, \cdots, x_{\lambda}, p_{v}^{i}\right)$. This message arrives $2 \lambda$ rounds later, i.e., in round $2 i-2$ at the swarm $S_{i-1}\left(x_{\lambda}\right)$. Then, within two rounds, Alg-LDS first constructs the handover graph $H_{i-1}$ and a new LDS $D_{i}$ based on these positions. This ensures that in the even rounds the forwarding step from Alg-Routing can be performed and the handover in the odd rounds. Thus, AlG-LDS maintains a routable overlay.

The algorithm proceeds in rounds. In any given round, every mature node picks a random position $p \in[0,1)$ and routes its $I D$ using message $\operatorname{Join}(I D, p)$ to the respective target points using Alg-Routing. This is the position the node will occupy in $D_{i}$ in round $2 i$.

It uses the routing algorithm Alg-Routing to send its ID to the swarms $S_{i-1}(p), S_{i-1}(p / 2)$, and $S_{i-1}(p+1 / 2)$ thereby creating the handover graph $H_{i-1}$ in round $2 i-1$. Starting from the handover graph $H_{i-1}$, the overlay $D_{i}$ can then be created in a single additional round through local introductions.

We will now describe the construction of the overlays $H_{i-1}$ and $D_{i}$ in detail given that the algorithm worked correctly until that round. We assume the system is currently in an even round $t=2 i-2$ and in all previous rounds the mature nodes formed a routable overlay $\mathcal{D}:=D_{0}, H_{0}, \ldots, D_{i-1}$. In other words, all messages of the from $\left(v, p_{v}^{i}\right)$ where $v$ is an identifier and $p_{v}^{i}$ is its position in $D_{i}$, are one round away from reaching their target. Then, the construction of $H_{i-1}$ and $D_{i}$ is as follows.

Listing 3: Overlay Maintenance Algorithm Alg-LDS


1. In round $2 i-2$ Alg-Routing executes the forwarding step. This implies, all messages are routed to their target location. In particular, each message with target point $p$ will be received by all nodes in $S_{i-1}(p)$. Recall that this is ensured by the fact that the message is sent to all nodes in $S_{i-1}(p)$ in the last step of the trajectory. Additionally, note that Alg-LDS ensures that the message is also forwarded to all nodes in $\left\langle p \pm \frac{2 c \lambda}{n}\right\rangle,\left\langle\frac{p}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle$, and $\left\langle\frac{p+1}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle$.
2. In round $2 i-1$, each node $w \in\left\langle p_{v}^{i} \pm \frac{2 c \lambda}{n}\right\rangle \cup\left\langle\frac{p_{v}^{i}}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle \cup\left\langle\frac{p_{v}^{i}+1}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle$ receives messages of the form $\operatorname{Join}\left(v, p_{v}^{i}\right)$. Therefore, the construction of the handover graph $H_{i-1}$ follows directly from the correctness of Alg-Routing and Definition 5. In the remainder of round $2 i-1$ two routines are processed concurrently. First, the handover step of Alg-Routing is executed. Second, all nodes must learn their neighbors in $D_{i}$ in order to execute the forwarding step in round $2 i$. For this, the nodes iterate over all received messages of the form $\left(v, p_{v}^{i}\right)$ and introduces them to all their neighbors. By introduction, we mean that the neighbor's $I D$ and position is sent to $v$. As we will see, for every pair of neighbors in $D_{i}$ there is at least one node that knows both their IDs and introduces them. These messages arrive in round $2 i$.
3. Finally, at the beginning of round $2 i$, each node knows all its neighbors in new overlay $D_{i}$ (and all messages that need to be forwarded in $D_{i}$ ). Thus, all mature nodes form the overlay $D_{i}$ (and can perform the forwarding step).

Note that after round $2 i+1$ the nodes' positions in $D_{i}$ and $D_{i+1}$ are in no relation with each other and hence, the edges in $D_{i+1}$ are independent of $D_{i}$. Therefore, the adversary stays oblivious of all nodes' current positions.

Further, observe that our approach requires that both the fresh and the mature nodes send out the join requests and that all messages take exactly the same time to reach its destination. The latter is ensured through Alg-Routing. For the former we assume that each fresh node is known by least one mature node, which is part of $D_{i}$. However, this will be ensured by Alg-Random and explained in the next section.

Listing 3 presents the pseudocode for the algorithm. Each node has the two variables $D_{t}^{u}$ and $H_{t}^{u}$. $D_{t}^{u}$ stores $u$ 's neighborhood in $D_{t}$ whereas $H_{t}^{u}$ stores the references for the handover. Both variables may be reset at the end of each round. The nodes pick a random position in the $[0,1)$-interval using a uniform hash function $h: V \times \mathbb{N} \rightarrow[0,1)$ known to all nodes. This hash function can either be established in the bootstrap phase by sending $O\left(\log ^{2} n\right)$ bits of shared randomness to each node or we can assume it is known to all nodes. The former case, it would have to be renewned every $O(\log (n))$ as the adversary could gain access to it. One way to do this would be to transform the routing algorithm into a broadcast algorithm that can send a given message to all nodes. For brevity we omit this and simply assume that all nodes know $h$ and the adversary does not have access to it.

This hash function takes the node's $I D$ and the current round as an input and computes a random value $p_{v}$. Note that this choice excludes some points in $[0,1)$ from being picked as all values need to be encoded in $O(\log n)$ bits. However, this does not impact the correctness of our algorithms. Instead, we just handle the values returned by $h$ as continuous values as it is a standard assumption (see, e.g. [14, 13] for the usage in overlay networks and 5] for a proof that these functions can indeed be simulated by few random bits).

Recall that the fresh nodes are not part of the overlay. Therefore, the mature nodes send out requests on behalf of each fresh node $u \in F_{t}$ known to them. Note that each node can compute $h(u, t)$ if it knows $u$ 's ID. The $I D$ s of these nodes are stored in the variable $C$. This variable is set by Alg-Random. Details on how it set can be found in the next section.

### 4.1 Analysis of Alg-LDS

In this section we show that Alg-LDS maintains a dynamic overlay with the properties needed for routing. Throughout this section we assume that Alg-Random works correctly and each fresh node is connected to $\Theta(\log n)$ mature nodes at any time. Thus, every node in the networks starts a join request in every even round.

Lemma 20 Let $\mathcal{D}_{t}:=\left(D_{0}, H_{0} \ldots, D_{i}\right)$ be routable graph until round $t=2 i$. Then it holds $\mathcal{D}_{t+2}:=$ $\left(D_{0}, H_{0} \ldots, D_{i}, H_{i}, D_{i+1}\right)$ is a routable graph until round $t+2$ w.h.p.

Proof: W.lo.g. we assume that the algorithm is currently in round $t=2 i$ and the overlays $\mathcal{D}_{t}:=$ $\left(D_{0}, H_{0} \ldots, D_{i}\right)$ were routable. This implies that the mature nodes know all neighbors in $D_{i}$, all join requests $\left(v, p_{v}^{i+1}\right)$ started $2 \lambda$ rounds ago are delivered, and the nodes are ready to perform the final forwarding step of the messages. Due to Lemma 13 we know that $D_{i}$ is good, and at least $3 / 4$-fraction of each swarm in $D_{i}$ survives until $2 i+1$. We will now show that Alg-LDS maintains the following three properties.

1. Alg-LDS successfully constructs $H_{i}$ in round $2 i+1$,
2. constructs a new LDS $D_{i+1}$ in round $2 i+2$, and
3. all swarms $S_{t+2}(p)$ are good w.h.p.

Together, these three properties imply that $\mathcal{D}_{t+2}:=\left(D_{0}, H_{0} \ldots, D_{i}, H_{i}, D_{i+1}\right)$ are a series of routable overlay.

The following lemma shows that Alg-LDS constructs $H_{i}$ in round $2 i+1$, i.e., we show that for any $p \in[0,1)$, every node in $S_{2 i}(p)$ knows the ID of every node in $S_{2 i+1}(p)$. The proof essentially follows using correctness of Alg-Routing and Lemma 5

Lemma 21 (Correctness of the Handover Construction) Let $\mathcal{D}_{t}:=\left(D_{1}, H_{1} \ldots, D_{i}\right)$ be routable graph until round $t=2 i$. Then, in round $2 i+1$, each node in $\left\langle p_{v}^{i+1} \pm \frac{2 c \lambda}{n}\right\rangle \cup\left\langle\frac{p_{v}^{i+1}}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle \cup\left\langle\frac{p_{v}^{i+1}+1}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle$ receives $\left(v, p_{v}^{i+1}\right)$ w.h.p. This implies that in round $2 i+1$ the nodes form the Handover graph $H_{i}$.

Proof: The proof follows directly using correctness of Alg-Routing and the overlay's topology. Since $\mathcal{D}_{t}$ is routable until (and including) round $2 i$, all messages that were started in round $2 i-2 \lambda$ are correctly routed to their target swarm w.h.p. via Alg-Routing. This includes all $\operatorname{Join}\left(v, p^{i+1}\right)$ messages that were started in round $2 i-2 \lambda$. Due to Lemma 13 we know that $D_{t}$ is good w.h.p. Thus, in round $2 i+1$ every node in $\left\langle p_{v}^{i+1} \pm \frac{2 c \lambda}{n}\right\rangle \cup\left\langle\frac{p_{v}^{i+1}}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle \cup\left\langle\frac{p_{v}^{i+1}+1}{2} \pm \frac{3 c \lambda}{2 n}\right\rangle$ received Join $\left(v, p^{i+1}\right)$ and therefore knows $v$ w.h.p. The proof then follows from the definition of the Handover graph $H_{i}$.

We continue with the construction of $D_{i+1}$. In particular, we show that every mature node $v \in$ $S_{i+1}\left(p_{v}^{i+1}\right)$ creates an edge to each of its new neighbors in $v^{\prime} \in S_{i+1}\left(p_{v}^{i+1}\right) \cup S_{i+1}\left(\frac{p_{v}^{i+1}}{2}\right) \cup S_{i+1}\left(\frac{p_{v}^{i+1}+1}{2}\right)$. We divide the neighbors into two sets.

1. the list neighbors left and right of $p_{v}^{i+1}$, and
2. the DeBruijn neighbors left and right of $\frac{p_{v}^{i+1}}{2}$ and $\frac{p_{v}^{i+1}+1}{2}$.

The following lemma show that for all nodes $v$ and $v^{\prime}$ which will be neighbors in $D_{i+1}$, w.h.p. there is at least one node $w$ that receives the messages $\left(v, p_{v}^{t+1}\right)$ and $\left(v^{\prime}, p_{v^{\prime}}^{t+1}\right)$ in round $2 i+1$ and thus introduces the nodes.

Lemma 22 Let $v, w$ be any two neighbors in $D_{i+1}$, then w.h.p.

$$
\mid\left\{u \in G_{2 i+1} \mid u \text { receives }\left(v, p_{v}^{i+1}\right) \text { and }\left(w, p_{w}^{i+1}\right)\right\} \left\lvert\, \geq \frac{14}{17} c \lambda\right.
$$

where $G_{2 i+1} \subseteq V_{2 i+1}$ are the set of good nodes in the round $2 i+1$.
Proof: Consider two nodes $v$ and $w$ with $d\left(p_{v}^{i+1}, p_{w}^{i+1}\right) \leq \frac{2 c \lambda}{n}$, i.e., neighbors in $D_{i+1}$. W.l.o.g. we assume that $p_{w}^{i+1}$ is right of $p_{v}^{i+1}$ and $d\left(p_{v}^{i+1}, p_{w}^{i+1}\right)=\frac{2 c \lambda}{n}$. We make these simplifying assumptions since (a) the proof is analogous for the left and right side and (b) any closer point can only have more nodes for the introduction.

Observe that the last step of Alg-Routing is executed in round 2i. Particularly, Alg-LDS ensures that the message $\operatorname{Join}\left(v, p_{v}^{i+1}\right)$ and $\operatorname{Join}\left(w, p_{w}^{i+1}\right)$ is forwards to every node in the interval $\left[p_{v}^{i+1}, p_{v}^{i+1}+\frac{2 c \lambda}{n}\right]$ and $\left[p_{w}^{i+1}, p_{w}^{i+1}-\frac{2 c \lambda}{n}\right]$, respectively. This is possible due to the topology of $D_{i}$. This implies there is a interval $I \in[0,1)$ of length $\frac{2 c \lambda}{n}$ such that all nodes belonging to this interval receive both $\operatorname{Join}\left(v, p_{v}^{i+1}\right)$ and $\operatorname{Join}\left(w, p_{w}^{i+1}\right)$ in round $2 i+1$. The claim then follows using Lemma 6 that interval $I$ has at least $c \lambda$ nodes w.h.p. and Lemma 13 that at least $\frac{14}{17}$ of those nodes in the interval $I$ are good nodes and remain until round $2 i+2$ w.h.p.

The next lemma shows that a $(2,2 \lambda+7)$-late adversary effectively reduces the adversarial churn to a randomized churn as the adversary is oblivious to which nodes belongs to which swarm in any given round.

Lemma $23 A(2,2 \lambda+5)$-late adversary enables Alg-LDS construct $D_{i+1}$ independent of $D_{i}$.
Proof: The proof follows from the correctness of Alg-Routing and Alg-LDS. Recall that in every even round, each mature node in $v \in M_{2 i-(2 \lambda+2)}$ picks a position $p_{v}^{i}$ in $D_{i}$ for itself and also for the fresh nodes that are connected to them. This is the position the node $v$ occupies in round $2 i$ i.e., $\operatorname{LDS} D_{i}$. $\operatorname{Join}\left(v, p_{v}^{i}\right)$ is routed using Alg-Routing and arrives at $p_{v}^{i}$ in round $2 i-1$ i.e., $H_{i-1}$ and then within a round Alg-LDS constructs $D_{i}$. Therefore, a $(2,2 \lambda+5)$-late adversary is oblivious to the position of node $v$ until round $2 i+2$. However, observe that ALg-LDS ensures that the node is at the position $p_{v}^{i+1}$ picked uniformly at random from $[0,1$ )-interval in round $2 i+2$. This implies that the position of node $v$ in $D_{i+1}$ is independent of $D_{i}$.

Finally, Lemma 20 follows from Lemma 21, 22, and 23. This concludes the analysis of the maintenance algorithm.

## Handling New and Fresh Nodes

We now present Alg-Random in detail. This algorithm ensures that each fresh node is known by $\delta \in \mathcal{O}(\log n)$ randomly chosen good mature nodes each round w.h.p. Algorithm is executed in rounds on nodes set $F_{t}$ and $M_{t}$ corresponding to the fresh and mature nodes of round $t$, respectively. Recall that every fresh node joins the network via a node which has been in the network for at least two rounds. This enables the bootstrapping node to update the newly joined node with IDs of $O(\log (n))$ mature nodes and also advertise the ID of the newly joined node to $O(\log (n))$ mature nodes in the overlay in the subsequent round. Each fresh node $f \in F_{t}$ which is at least one round old, advertises its own ID to $O(\log (n))$ mature nodes in the overlay. Every such unique advertisement a mature node receives, is associated with a unique key in $[0, O(\log (n))]$ and stored in its memory. Each mature node $m \in M_{t}$, uniformly and independently at random samples $O(\log (n))$ other mature nodes in the overlay using ALG-SAMPLING. Each sampled ID of a matured node is either sent to a newly joined node (i.e. less than a round old and bootstrapped via $m$ ) with probability $p=1 / 2$ or is with probability $1-p$ forwarded to the ID of a fresh node, if available, whose key is picked uniformly at random from $[0, O(\log (n))]$.

Note that at the end of round $t$ a node forgets all its incoming connections from fresh nodes and the assignment of numbers to $I D$ s is reset. Last, note that the bootstrap phase ends once the first tokens reach their target.

Listing 4 depicts the pseudocode for Alg-Random. We use two types of messages, TOKEN(v) and CONNECT(v). Both messages only contain a nodes $v$ 's $I D$. The former is used to spread the mature nodes' $I D \mathrm{~s}$, the latter is used send a fresh node's $I D$ for sampling. Note that all token that are ready to create an are stored in the variable $T$. Further, the array $\left(c_{1}, \ldots, c_{2 \delta}\right)$ stores the assignment of numbers to $I D \mathrm{~s}$. It holds $c_{i}=v$ if $v$ 's $I D$ is assigned to $i$. If no $I D$ is assigned to $i$ we write $c_{i}=\perp$. Note that the set $C$ mentioned in Listing 3 consists of all $c_{i} \neq \perp$. Last, note that a node can distinguish whether it received a $\operatorname{TOKEN}(v)$ message through Algorithm Alg-Sampling, i.e., in step 1 of the sampling process sketched above, or directly from a node, i.e., in step 3.

Listing 4: Random Overlay Algorithm Alg-Random

```
1 Desc: In each round \(t\) each fresh node connects to \(\delta\) mature nodes that joined at
    least \(t-(2 \lambda+5)\) rounds ago.
2
3 Note: The following code is executed by each node \(u \in V\) every round \(t\). All types
        of messages are received in the given order. The last block of commands is
        executed after all messages have been handled.
4
5 Round \(t\)
6 Upon receiving \(T O K E N(v)\) from node \(u^{\prime}\) :
\(7 \quad T \longleftarrow T \cup\{v\} \quad \triangleright\) Tokens ready to be used
8
9 Upon receiving \(C O N N E C T(v)\) from node \(v\) :
if \(\exists i \in[0,2 \delta]\) with \(c_{i}=\perp\)
\(11 \quad i \longleftarrow\) number chosen uniformly from all \(i \in[0,2 \delta]\) with \(c_{i}=\perp\)
\(12 \quad c_{i} \longleftarrow v\)
13
14 Upon \(v\) joining:
\(15\left(w_{1}, \ldots, w_{\delta}\right) \longleftarrow \delta\) tokens chosen u.i.r from \(T\)
16 Send \(C O N N E C T(v)\) to all \(w_{1}, \ldots, w_{\delta} \quad \Delta u\) sends on behalf of \(v\)
\(17\left(w_{1}, \ldots, w_{\delta}\right) \longleftarrow \delta\) tokens chosen u.i.r from \(T\)
18 Send \(\operatorname{TOKEN}\left(w_{1}\right), \ldots, \operatorname{TOKEN}\left(w_{\delta}\right)\) to \(v \quad \triangleright\) Supply \(v\) with tokens
19
20 Upon receiving \(T O K E N(v)\) through Algorithm Alg-Sampling:
\(x \longleftarrow\) uniformly chosen from \(\{0,1\}\)
    if \(x=0\) :
        \(c_{i} \longleftarrow\) random element from \(c_{i}, \ldots, c_{2 \delta}\)
        Send TOKEN \((v)\) to \(c_{i}\) (or discard if \(c_{i}=\perp\) )
    else if \(u\) :
        \(T \longleftarrow T \cup\{v\} \quad \triangleright\) Tokens ready to be used
27
28 Finally:
    if \(u\) is fresh:
        \(\left(w_{1}, \ldots, w_{\delta}\right) \longleftarrow \delta\) tokens chosen u.i.r from \(T\)
        Send \(C O N N E C T(u)\) to all \(w_{1}, \ldots, w_{\delta}\)
    else if \(u\) is mature:
        Send TOKEN(v) to \(\tau\) random nodes using Algorithm Alg-SAmpling
    \(\left(c_{1}, \ldots, c_{2 \delta}\right) \longleftarrow(\perp, \ldots, \perp) \quad\) Reset IDs
    \(T \longleftarrow \emptyset \quad \triangleright\) Drop unused tokens
```


### 4.2 Analysis of Alg-Random

In this section we show that every fresh node is able is send its $I D$ to $\delta$ mature nodes each round w.h.p. and thus stays connected to the network. In particular, we assume that $\delta \in O(\log n)$. Therefore, we prove the following lemma.

Lemma 24 (Random Overlay Lemma) Assume that until round $t-1$ each fresh node was connected to at least $\frac{1}{2} \delta$ good nodes each round. Then, it holds w.h.p. that each $v \in F_{t}$ successfully connects to $\frac{1}{2} \delta$ good nodes.
We prove the lemma in several steps. First, we show that each node receives $\Omega(\tau)$ tokens w.h.p. To prove this we make use of a simple balls-into-bins argument. Recall that each mature node in starts $\tau$ tokens in round $t-(2 \lambda+5)$ that reach their random destination in round $t$. Further, we can show that the tokens are uniformly distributed among all nodes.

Lemma 25 Assume Lemma 24 held until round $t-1$. Further, let $X(\theta, v)$ denote the event that any token $\theta$ reaches $v$ in round $t$. Then the following statements hold:

1. Any token reaches $v \in V_{t}$ with the same probability, i.e.,

$$
\operatorname{Pr}[X(\theta, v)=1]=\operatorname{Pr}\left[X\left(\theta^{\prime}, v\right)=1\right] .
$$

2. For each token $\theta$, it holds

$$
\operatorname{Pr}[X(\theta, v)=1] \geq \frac{1}{32 n} .
$$

## Proof:

1. We extend Lemma 18 to fresh nodes and show all token reach a node $v \in V_{t}$ with the same (but not necessarily uniform) probability. Consider a token $\theta$ independently of its source node and let $X(\theta, u)$ indicate that $\theta$ reaches $u$. Further, let $V_{t-1}^{\prime} \subset M_{t-1}$ be the set of all nodes that know $v$ 's $I D$. If $v$ receives $\theta$ in round $t$, then the following two events must happen
(a) The token must be sent to any mature node $v^{\prime} \in V_{t-1}$ using Alg-Sampling. We denote this event as $X_{1}\left(\theta, v^{\prime}\right)$.
(b) Given any $v^{\prime} \in V_{t-1}^{\prime}$ received $\theta$, it must forwarded is to $v$ in round $t-1$. We denote this event as $X_{2}^{v^{\prime}}(\theta, v)$.

We can easily show that both these events have the same probability for two tokens of possibly different origin. The uniformity of the first event directly follows from Lemma 18 . Here, we showed that $\operatorname{Pr}\left[X_{1}\left(w, v^{\prime}\right)\right]=\operatorname{Pr}\left[X_{1}\left(u, v^{\prime}\right)\right]$ for every $u, w \in V_{t-(2 \lambda+5)}$ and $v \in V_{t}$. Note that for two different $v^{\prime}, v^{\prime \prime} \in V_{t}$ the probabilities $\operatorname{Pr}\left[X_{1}\left(w, v^{\prime}\right)\right]$ and $\operatorname{Pr}\left[X_{1}\left(w, v^{\prime}\right)\right]$ may differ. The uniformity of the second event follows from the fact that each token is forwarded to $v$ with probability of exactly $\frac{1}{4 \delta}$. To finalize the proof, consider two nodes $u, w \in V_{t-\lambda}$ and let $\theta^{u}$ and $\theta^{w}$ be tokens sent by $v$ and $w$ respectively. Then,

$$
\begin{aligned}
\operatorname{Pr}\left[X\left(\theta^{u}, v\right)=1\right] & =\sum_{v^{\prime} \in V_{t-1}^{\prime}} \operatorname{Pr}\left[X_{1}\left(\theta^{u}, v^{\prime}\right)=1 \cap X_{2}^{v^{\prime}}\left(\theta^{u}, v\right)=1\right] \\
& =\sum_{v^{\prime} \in V_{t-1}^{\prime}} \operatorname{Pr}\left[X_{1}\left(\theta^{u}, v^{\prime}\right)=1\right] \cdot \operatorname{Pr}\left[X_{2}^{v^{\prime}}\left(\theta^{u}, v\right)=1 \mid X_{1}\left(\theta^{u}, v^{\prime}\right)=1\right] \\
& =\sum_{v^{\prime} \in V_{t-1}^{\prime}} \operatorname{Pr}\left[X_{1}\left(\theta^{w}, v^{\prime}\right)=1\right] \cdot \operatorname{Pr}\left[X_{2}^{v^{\prime}}\left(\theta^{w}, v\right)=1 \mid X_{1}\left(\theta^{w}, v^{\prime}\right)=1\right] \\
& =\operatorname{Pr}\left[X\left(\theta^{w}, v\right)=1\right] .
\end{aligned}
$$

Here, the first equality is due to the law of total probability and second equality is due to Lemma 18 and the fact that each mature node $v^{\prime}$ forwards a token to a fresh node with probability $\frac{1}{4 \delta}$.
2. The fact that $\operatorname{Pr}[X(v, w)=1] \in \Omega\left(\frac{1}{n}\right)$ then follows from three facts:
(a) First, a token reaches a given mature node with probability at least $\frac{1}{4 n}$. This follows directly from Lemma 18 ,
(b) Second, each fresh node is connected to at least $\frac{\delta}{2}$ mature nodes w.h.p. This follows because we assume that Lemma 24 holds true in round $t-1$.
(c) Last, a mature node forwards a token to a connected node with probability $\frac{1}{4 \delta}$.

Combining these three facts yields the result. Formally:

$$
\begin{aligned}
\operatorname{Pr}\left[X(\theta, v)=1| | V_{t-1}^{\prime} \mid \geq \delta / 2\right] & =\sum_{v^{\prime} \in V_{t-1}^{\prime}} \operatorname{Pr}\left[X_{1}\left(\theta, v^{\prime}\right)=1\right] \cdot \operatorname{Pr}\left[X_{2}^{v^{\prime}}(\theta, v)=1 \mid X_{1}\left(\theta, v^{\prime}\right)=1\right] \\
& \geq \sum_{v^{\prime} \in V_{t-1}^{\prime}} \frac{1}{4 n} \frac{1}{4 \delta} \\
& \geq \frac{\delta}{2} \frac{1}{4 n} \frac{1}{4 \delta} \\
& \geq \frac{1}{32 n}
\end{aligned}
$$

Lemma 26 Assume Lemma 24 held until round $t-1$. Further, let $X(u, v)$ denote the event that any token sent by $u$ reaches $v$ in round $t$. Then the following statements hold:

1. Each node sends at least one token to node $v \in V_{t}$ with the same probability, i.e., $\forall u, w \in V_{t-(2 \lambda+5)}$

$$
\operatorname{Pr}[X(u, v)=1]=\operatorname{Pr}[X(w, v)=1]
$$

2. For each $u \in V_{t-(2 \lambda+2)}$ it holds

$$
\operatorname{Pr}[X(u, v)=1] \geq \frac{\tau}{33 n}
$$

## Proof:

1. Now, we observe the variables $X(u, v)$ and $X(w, v)$ that denotes if any of $u$ 's or $w$ 's tokens reached $v$. Recall that both $u$ and $w$ send $\tau$ tokens. We denote these tokes as $\theta_{1}^{u}, \ldots, \theta_{\tau}^{u}$ and $\theta_{1}^{u}, \ldots, \theta_{\tau}^{u}$. Let $X(\theta, v)$ be defined as in Lemma 25. The probability that any of these tokens reach $v$ is given by:

$$
\operatorname{Pr}\left[\bigcup_{i=1, \ldots, \tau} X\left(\theta_{i}^{u}, v\right)=1\right]=1-\operatorname{Pr}\left[\bigcap_{i=1, \ldots, \tau} X\left(\theta_{i}^{u}, v\right)=0\right]
$$

Since all these tokens are independent, it holds that:

$$
\operatorname{Pr}\left[\bigcap_{i=1, \ldots, \tau} X\left(\theta_{i}^{u}, v\right)=0\right]=\prod_{i=1, \ldots, \tau}\left(1-\operatorname{Pr}\left[X\left(\theta_{i}^{u}, v\right)=1\right]\right)
$$

The same holds respectively for $\operatorname{Pr}\left[\bigcap_{i=1, \ldots, \tau} X\left(\theta_{i}^{w}, v\right)=0\right]$. Putting these observations together, we get that:

$$
\operatorname{Pr}[X(u, v)=1]=\operatorname{Pr}\left[\bigcup_{i=1, \ldots, \tau} X\left(\theta_{i}^{u}, v\right)=1\right]=1-\operatorname{Pr}\left[\bigcap_{i=1, \ldots, \tau} X\left(\theta_{i}^{u}, v\right)=0\right]
$$

$$
\begin{aligned}
& =1-\prod_{i=1, \ldots, \tau}\left(1-\operatorname{Pr}\left[X\left(\theta_{i}^{u}, v\right)=1\right]\right) \\
& =1-\prod_{i=1, \ldots, \tau}\left(1-\operatorname{Pr}\left[X\left(\theta_{i}^{w}, v\right)=1\right]\right) \\
& =1-\operatorname{Pr}\left[\bigcap_{i=1, \ldots, \tau} X\left(\theta_{i}^{w}, v\right)=0\right]=\operatorname{Pr}\left[\bigcup_{i=1, \ldots, \tau} X\left(\theta_{i}^{w}, v\right)=1\right] \\
& =\operatorname{Pr}[X(w, v)=1] .
\end{aligned}
$$

This was to be shown.
2. For any pair of $v$ and $u$, the probability is lower bounded by

$$
\begin{aligned}
\operatorname{Pr}[X(u, v)=1] & =1-\operatorname{Pr}[X(u, v)=0] \\
& =1-\prod_{i=1, \ldots, \tau}\left(1-\operatorname{Pr}\left[X\left(\theta_{i}^{u}, v\right)=1\right]\right) \\
& =1-\left(1-\frac{1}{32 n}\right)^{\tau} \\
& \geq 1-\exp \left(-\frac{\tau}{32 n}\right) \\
& \geq 1-\left(1-\frac{\tau}{32 n}+\left(\frac{\tau}{32 n}\right)^{2}\right) \\
& \geq \frac{\tau}{32 n}-\left(\frac{\tau}{32 n}\right)^{2} \\
& \geq \frac{\tau}{33 n}
\end{aligned}
$$

where for the second inequality we use the fact that for all $x \leq 1$,

$$
\exp (x) \leq 1+x+x^{2}
$$

The last inequality follows from fact that $\tau \in O(\log n)$ and thus $\frac{\tau}{32 n}$ can be made arbitrarily small for a big enough $n$.

Together with our assumptions on the churn rate, we get the following lemma:
Lemma 27 Let $c \geq 280 k$. Let each mature node start $\tau \geq 20 c \lambda$ token, then each fresh node receives at least $\frac{\tau}{100}$ distinct token with probability at least $1-\frac{1}{n^{k}}$.
Proof: Recall that at least $\frac{15}{16} n$ mature nodes in round $t-(2 \lambda+5)$ that start $\tau$ tokens each. Hence, the minimal number of nodes that start tokens is at least $K:=\frac{15 n}{16}$. Further, there are at most $\frac{17}{16} n$ nodes in round $t$. Fix a node $v$ and let $X_{1}, \ldots, X_{K}$ be the indicator variables that a nodes has a token that reaches $v$. Then the expected number of distinct tokens received by node $v$ is given by,

$$
E\left[\sum_{i=1}^{K} X_{i}\right] \geq \sum_{i=1}^{K} \frac{\tau}{33 n} \geq \frac{15 \tau n}{16 \cdot 33 \cdot n} \geq \frac{\tau}{50}
$$

where we use $\operatorname{Pr}\left[X_{i}=1\right]=\frac{\tau}{33 n}$ due to Lemma 26. Given that all mature nodes send their tokens independent of one another, the Chernoff Bound is applicable and the lemma follows for a big enough c. In particular, it holds for $c \geq 280 k$ :

$$
\operatorname{Pr}\left[X \leq \frac{\tau}{100}\right]=\operatorname{Pr}\left[X \leq(1-1 / 2) \frac{\tau}{50}\right]
$$

$$
\begin{aligned}
& \leq \exp \left(-\frac{\tau}{4 \cdot 50 \cdot 3}\right) \\
& \leq \exp (-k \lambda)=n^{-k}
\end{aligned}
$$

This basically tells us that - as long as we choose $\tau$ bigger than $20 c \lambda$ - each node will receive roughly $\Omega(\tau)$ distinct tokens w.h.p, which it can then use to advertise itself and the new nodes connected to it.

Next, we need to consider, how big we need to choose $\tau$ such that each node has enough tokens to ensure that it is able to connect to $\frac{\delta}{2}$ mature nodes. Lemma 27 gives us that choosing $\tau$ such that $\frac{\tau}{100} \geq(\phi+1) \cdot \delta$, where $\phi=O(1)$ is the maximum number of nodes that could join via a fresh node in any given round, then each fresh node has enough tokens to advertise itself to $\delta$ distinct mature nodes every round and also provide for the newly joined nodes.

In the following we can assume that each fresh node sends a connection request to $\delta$ nodes. However, these requests can still fail for two reasons:

1. First, the $I D$ of the token used for the connections belongs to a node that has been churned out.
2. Second, the the target has received more than $2 \delta$ connection requests and refuses the connection.

The first factor depends on the number of nodes have been churned out and on the numbers of connections we make. The second term only depends on the random process that creates these edges.

We begin by showing that only a small fraction of connection request are sent to churned out nodes. In the following lemma we say node $v$ is good in round $t$ if and only if $v \in V_{t-(2 \lambda+5)} \cap V_{t+2}$ and referred to as bad, otherwise.
Lemma 28 Suppose that $\tau \geq 26000 k \lambda$ and $\delta \geq 60 k \lambda$, then each fresh node has at least $\frac{\delta}{2}$ connections to good nodes with probability at least $1-\frac{1}{n^{k}}$.
Proof: The proof of this lemma is straightforward and mostly technical. The basic outline is as follows: Due to its lateness the adversary cannot anticipate where a node will send its tokens. Thus, the tokens of good and bad nodes will randomly spread to the fresh nodes. As will see, in expectation each fresh node roughly receives a $\frac{15}{16}$ fraction of good nodes. Since the sampling is independent, this implies there is at least a $\frac{13}{18}$ fraction w.h.p (for a big enough $\tau$ ) due to the Chernoff Bound. Since a fresh node randomly draws its connection without replacement there are also $\frac{13}{18} \delta$ successful connections in expectation. Since drawing without replacement is NA, another application of the Chernoff bound concludes the proof.

We will now prove these claims in more detail: Fix a node $v \in V_{t}$ that advertises itself or a newly joined node. Let $X_{1}, \ldots, X_{\delta}$ be the binary RVs such that $X_{j}$ denotes if $j^{\text {th }}$ advertisement by $v$ is successful, i.e., its identifier is advertised to a good node. The outcome of $X=\sum_{i} X_{i}$ depends on two values, the overall number of tokens that $v$ received and the number of identifiers of good nodes.

We have already established that the number of distinct tokens that a node receives can be subjected to the Chernoff Bound and is therefore concentrated around its expectation. The same holds for the number of good identifiers. Let $G$ be the number of good identifiers that $v$ draws from the set of available tokens. One can easily verify that $G$ is the sum independent binary random variables: For each good node $w \in V_{t-(2 \lambda+5) \cap V_{t+2}}$, let $G_{w} \in\{0,1\}$ be the indicator for the event that $w$ sends one token with its identifier to $v$. Then it holds $G:=\sum_{w \in V_{t-(2 \lambda+5)} \cap V_{t+2}} G_{w}$ and all $G_{w}$ 's are independent.

Recall that at least $n\left(1-\frac{1}{16}\right)$ and at most $n\left(1+\frac{1}{16}\right)$ nodes started tokens $2 \lambda+5$ rounds ago. Since at most a $1 / 16$-fraction of all nodes that started a token are churned out until round $t+2$, it holds that $\left(\frac{15}{16}\right)^{2} n$ is a lower bound for the number of good tokens started $2 \lambda+5$ rounds ago.

Let now $Y$ be the number of all distinct tokens that $v$ received. Let $p \in\left[\frac{\tau}{33 n}, \frac{\tau}{n}\right]$ be the probability that at least token of a fixed node reaches $v$. Then, it holds $E[Y] \leq p n\left(1+\frac{1}{16}\right)$ and $E[G] \geq p\left(\frac{15}{16}\right)^{2} n$. We will upper and lower bound $Y$ and $G$ respectively. We start with $Y$. Assuming that $p n \geq 768 k \lambda$ the Chernoff Bound gives us that,

$$
\begin{equation*}
\operatorname{Pr}\left[Y \geq \frac{9}{8} p n\right]=\operatorname{Pr}[Y \geq(1+1 / 17)(1+1 / 16) p n] \leq \exp \left(-\frac{p n}{2 \cdot 17 \cdot 16}\right) \leq \exp (-k \lambda)=n^{-k} \tag{9}
\end{equation*}
$$

$$
\begin{equation*}
\operatorname{Pr}\left[G \leq \frac{13}{16} p n\right] \leq \operatorname{Pr}\left[G \leq(1-1 / 15)\left(\frac{15}{16}\right)^{2} p n\right] \leq \exp \left(-\frac{15^{2} \cdot p n}{3 \cdot 16^{2} \cdot 15^{2}}\right) \leq \exp (-k \lambda)=n^{-k} \tag{10}
\end{equation*}
$$

Therefore, it remains to show that we can choose $p n$ big enough for these statements to hold. Recall that $\frac{\tau}{n} \geq p \geq \frac{\tau}{33 n}$, then for $\tau \geq 26000 k \lambda$, we have that $p n \geq 768 k \lambda$.

Now, we condition on (9) and (10) being false and denote this event as $\mathcal{G}$. In this case, a simple calculation reveals that at least a $\frac{13}{18}$-fraction of tokens is good. If we pick $\delta^{\prime}=\min \{\delta, Y\}$ of these tokens uniformly at random without replacement, a constant fraction will point to good nodes in expectation:

$$
\begin{aligned}
E[X \mid \mathcal{G}] & \geq \sum_{g=\frac{13}{16} p n}^{n\left(1+\frac{1}{16}\right)} \sum_{y=0}^{\frac{9}{8} p n} \operatorname{Pr}[G=g, Y=y \mid \mathcal{G}] \cdot \delta^{\prime} \frac{g}{y} \\
& \geq \sum_{g=\frac{13}{16} p n}^{n\left(1+\frac{1}{16}\right)} \sum_{y=0}^{\frac{9}{8} p n} \operatorname{Pr}[G=g, Y=y \mid \mathcal{G}] \cdot \delta^{\prime} \frac{\frac{13}{16} p n}{\frac{9}{8} p n} \\
& =\delta^{\prime} \frac{\frac{13}{16} p n}{\frac{9}{8} p n} \sum_{g=\frac{13}{16} p n}^{n\left(1+\frac{1}{16}\right)} \sum_{y=0}^{\frac{9}{8} p n} \operatorname{Pr}[G=g, Y=y \mid \mathcal{G}] \\
& =\delta^{\frac{13}{\frac{1}{16}} p n} \frac{\frac{9}{8} p n}{} \\
& =\delta^{\prime} \frac{13}{18} .
\end{aligned}
$$

We show that under these circumstances at least half of all advertisements go to good nodes for a big enough $\delta^{\prime}$. Note that we observe a hyper-geometric distribution, which is known to be NA (cf. [7]). Thus, by the Chernoff Bound, a constant fraction points to living nodes w.h.p. if we choose $\delta^{\prime}$ high enough. In particular, by choosing $\delta \geq 60 k \lambda$

$$
\begin{array}{rlr}
\operatorname{Pr}\left[\left.\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta \right\rvert\, \mathcal{G}\right] & \leq \operatorname{Pr}\left[\sum_{i=0}^{\delta} X_{i} \leq(1-5 / 18) \frac{13}{18} \delta\right] \\
& \leq \exp \left(-\frac{5^{2} \cdot 13 \delta}{3 \cdot 18^{3}}\right) \leq \exp \left(-\frac{\delta}{60}\right) \quad \triangleright U \operatorname{sing} \frac{5^{2} \cdot 13}{3 \cdot 18^{3} \geq \frac{1}{60}} \\
& \leq \exp (-k \lambda)=n^{-k} . & \triangleright \text { for } \delta \geq 60 k \lambda
\end{array}
$$

Note that since $\mathcal{G}$ holds w.h.p, we have that

$$
\begin{aligned}
\operatorname{Pr}\left[\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta\right] & =\operatorname{Pr}[\mathcal{G}] \cdot \operatorname{Pr}\left[\left.\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta \right\rvert\, \mathcal{G}\right]+\operatorname{Pr}[\neg \mathcal{G}] \cdot \operatorname{Pr}\left[\left.\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta \right\rvert\, \neg \mathcal{G}\right] \\
& \leq \operatorname{Pr}[\mathcal{G}] \cdot \operatorname{Pr}\left[\left.\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta \right\rvert\, \mathcal{G}\right]+\operatorname{Pr}[\neg \mathcal{G}] \cdot 1 \\
& =\operatorname{Pr}[\mathcal{G}] \cdot \operatorname{Pr}\left[\left.\sum_{i=0}^{\delta} X_{i} \leq \frac{1}{2} \delta \right\rvert\, \mathcal{G}\right]+\operatorname{Pr}\left[(Y \geq(1+1 / 15) p n) \cup\left(|G| \leq \frac{14}{16} p n\right)\right] \\
& \leq\left(1-\frac{1}{n^{k}}\right) \frac{1}{n^{k}}+\frac{2}{n^{k}} \leq \frac{3}{n^{k}} .
\end{aligned}
$$

Thus, the statement holds w.h.p. Finally, note that our choice of $\tau$ already implies that we receive that least $\delta$ tokens.

Lemma 29 A $(2,2 \lambda+7)$-late adversary enables Alg-Random ensure that every fresh node is connected to $\frac{\delta}{2}$ mature nodes in each round.
Proof: The proof follows using the correctness of Alg-Routing and Alg-Random. Note that the adversary is oblivious of the random edges because they only persist for 2 rounds. Each mature node disseminates tokens to random positions in the $[0,1)$ interval. The tokens arrive at their target node for being sampled after $2 \lambda+2$ rounds. The mature nodes that receive the token forward them to fresh nodes which in turn connect to the mature nodes to stay connected in the network until they mature themselves. The fresh nodes then receive new tokens from these connections. The entire process takes $2 \lambda+5$ rounds in total. Therefore, in any given round $t$ a $(2,2 \lambda+7)$ adversary is oblivious to any communication between the fresh nodes and mature nodes, since all connections established until round $t$ are already defunct, i.e., the adversary is unable to anticipate which tokens reach a given fresh node. This in turn enables Alg-Random maintain the invariant every round.

Lemma 29 and the churn parameters ensure that there exists a constant size set of good nodes $G_{t}:=V_{t-(2 \lambda+5)} \cap V_{t+2}$ that send a token in round $t-(2 \lambda+5)$ and are not churned out until round $V_{t+2}$. Therefore, if a node receives enough tokens of good nodes, it can successfully advertise its identifier w.h.p.

It remains to show that at most $2 \delta$ fresh nodes connect to a mature node in any given round. We first analyze the expected number of incoming connections.

Lemma 30 Let $A(u, v)$ denote the event that $v$ advertises itself to $u$. Then, $\forall u, w \in V_{t}$ and any two tokens $\theta, \theta^{\prime}$ it holds that,

$$
\operatorname{Pr}[A(u, v)]=\operatorname{Pr}[A(w, v)]
$$

Proof: We again divide the process into two stages. First, the token of $u$ needs to reach $v$ and then needs to be picked for the advertisement. Let $Y_{1}(u, v)$ and $Y_{1}(w, v)$ the respective events that tokens of $u$ and $w$ reached $v$. By Lemma 25 we already know that these events have the same probability.

Note that the actual choice of the nodes which are advertised only depends on the number of distinct available tokens. In particular, given that a node received $\ell$ distinct tokens, the probability for one of these tokens to be used is $\min \left\{1, \frac{\delta}{\ell}\right\}$. This follows from the fact that we draw (up to) $\delta$ tokens uniformly at random or all tokens if we received less than $\delta$. Thus, we draw without replacements and observe a hyper geometric distribution.

Let now $N_{v}$ be the number of distinct tokens received by $v$. Since all nodes send (at least) one token to $v$ independently and with same probability $p:=\operatorname{Pr}\left[Y_{1}(u, v)=1\right]$, the value of $N_{v}$ only depends on the number of nodes we observe. More formally, it holds:

$$
\begin{aligned}
\operatorname{Pr}\left[N_{v}=\ell\right] & :=\sum_{S \subset M_{t},|S|=\ell} \operatorname{Pr}\left[\bigcap_{u \in S} Y_{1}(u, v)=1\right] \cdot \operatorname{Pr}\left[\bigcap_{u^{\prime} \notin S} Y_{1}\left(u^{\prime}, v\right)=0\right] \\
& :=\sum_{S \subset M_{t},|S|=\ell} \prod_{u \in S} \operatorname{Pr}\left[Y_{1}(u, v)=1\right] \cdot \prod_{u^{\prime} \notin S} \operatorname{Pr}\left[Y_{1}\left(u^{\prime}, v\right)=0\right] \\
& =\binom{\left|M_{t}\right|}{\ell} p^{\ell}(1-p)^{n-\ell} .
\end{aligned}
$$

Thus, if we condition on the fact that $v$ already received a token of a certain node, the probability that this node receives $\ell-1$ additional tokens from different nodes stays the same. It holds,

$$
\begin{aligned}
\operatorname{Pr}\left[N_{v}=\ell \mid Y_{1}(u, v)=1\right] & =\sum_{S \subset M_{t} \backslash\{u\}:|S|=\ell-1} \operatorname{Pr}\left[\bigcap_{u^{\prime} \in S} Y_{1}\left(u^{\prime}, v\right)=1\right] \cdot \operatorname{Pr}\left[\bigcap_{u^{\prime \prime} \notin S} Y_{1}\left(u^{\prime \prime}, v\right)=0\right] \\
& =\binom{\left|M_{t}\right|-1}{\ell-1} \cdot p^{\ell-1} \cdot(1-p)^{\left|M_{t}\right|-(\ell-1)}
\end{aligned}
$$

$$
\begin{aligned}
& =\sum_{S \subset V \backslash\{w\}:|S|=\ell-1} \operatorname{Pr}\left[\bigcap_{u^{\prime} \in S} Y_{1}\left(u^{\prime}, v\right)=1\right] \cdot \operatorname{Pr}\left[\bigcap_{u^{\prime \prime} \notin S} Y_{1}\left(u^{\prime \prime}, v\right)=0\right] \\
& =\operatorname{Pr}\left[N_{v}=\ell \mid Y_{1}(w, v)=1\right] .
\end{aligned}
$$

Thus, when summing over all different outcomes, we get,

$$
\begin{aligned}
\operatorname{Pr}[A(u, v)=1] & =\operatorname{Pr}\left[Y_{1}(u, v)=1\right] \cdot \operatorname{Pr}\left[A(u, v) \mid Y_{1}(u, v)=1\right] \\
& =\operatorname{Pr}\left[Y_{1}(u, v)=1\right] \cdot\left(\sum_{\ell=1}^{\left|M_{t}\right|} \operatorname{Pr}\left[N_{v}=\ell \mid Y_{1}(u, v)\right] \cdot \min \left\{1, \frac{\delta}{\ell}\right\}\right) \\
& =\operatorname{Pr}\left[Y_{1}(w, v)=1\right] \cdot\left(\sum_{\ell=1}^{\left|M_{t}\right|} \operatorname{Pr}\left[N_{v}=\ell \mid Y_{1}(w, v)=1\right] \cdot \min \left\{1, \frac{\delta}{\ell}\right\}\right) \\
& =\operatorname{Pr}\left[Y_{1}(w, v)=1\right] \cdot \operatorname{Pr}\left[A(w, v) \mid Y_{1}(w, v)=1\right] \\
& =\operatorname{Pr}[A(w, v)=1] .
\end{aligned}
$$

This was to be shown.
Intuitively, this lemma implies that all node have the same probability of receiving an advertisement. Thus, since there are at $\delta n$ advertisement, the expected number of incoming connections is bound by $\delta$.

Lemma 31 Fix a mature node $w$ that started tokens $2 \lambda+5$ rounds ago. Let $X$ be a random variable that denotes the number of fresh nodes that advertise themselves to $w$. It holds:

$$
\mathbf{E}[X] \leq \delta
$$

Proof: Let $f:=\left|F_{t}\right|$ be number of fresh nodes and $m=\left|M_{t}\right|$ the number of mature nodes in round $t$. Note that $\frac{f}{m}$ is at most $\frac{1}{8}$ due to our choice of $\alpha$ and $\kappa$.

For a fixed $v \in F_{t}$ and $w \in M_{t}$ let $A(v, w)$ be the binary RV that denotes if $v$ connects to $w$. Let now $A_{v}:=\sum_{w \in M_{t}} A(v, w)$ be the random variable that counts $v$ 's advertisements. Since each node creates at most $\delta$ advertisements, it must hold that:

$$
\mathbf{E}\left[A_{v}\right] \leq \delta
$$

Further, we have that

$$
\mathbf{E}\left[A_{v}\right]=\sum_{w \in V} E[A(v, w)]
$$

This follows from the linearity of expectation. Given that $A(v, w)$ is a binary random variable, we also have that

$$
\mathbf{E}[A(v, w)]=\operatorname{Pr}[A(v, w)]
$$

Using Lemma 30 we also know that $\operatorname{Pr}[A(u, v)]=\operatorname{Pr}[A(w, v)]$. Now we combine our observations to bound $\operatorname{Pr}[A(u, v)]$. First, see that:

$$
\begin{array}{rlr}
\delta \geq \mathbf{E}\left[A_{v}\right]=\sum_{w \in V} \mathbf{E}[A(v, w)] & \triangleright \text { By linearity of Expectation } \\
& =\sum_{w \in V} \operatorname{Pr}[A(v, w)] & \triangleright A s \mathbf{E}[A(v, w)]=\operatorname{Pr}[A(v, w)] \\
& =m \cdot \operatorname{Pr}[A(v, w)] & \triangleright A s \operatorname{Pr}[A(v, w)]=\operatorname{Pr}\left[A\left(v, w^{\prime}\right)\right]
\end{array}
$$

Therefore,

$$
\operatorname{Pr}[A(v, w)] \leq \frac{\delta}{m}
$$

Let $X$ be a random variable that counts the number of incoming connections to the node $w$. Then,

$$
\mathbf{E}[X]=\sum_{v \in F_{t}} \mathbf{E}[A(v, w)] \leq \frac{\delta}{8}
$$

Thus, our claim holds.
Lemma 32 Each mature node receives at most $2 \delta$ connections from fresh nodes w.h.p.
Proof: Fix a mature node $w \in M_{t}$ and let $X_{v} \in\{0,1\}$ be the random variable that indicates whether $v \in F_{t}$ connects to $w$. Further, let $X:=\sum_{v \in F_{t}} X_{v}$ the sum of all fresh nodes that connect to $w$.

We will show that $\left(X_{v}\right)_{v \in F_{t}}$ is negatively correlated and then use the Chernoff Bound on $X:=$ $\sum_{v \in F_{t}} X_{v}$ to prove the lemma.

Let $A$ and $B$ be two disjoint subsets of fresh nodes and $X_{A}$ and $X_{B}$ the respective subsets of variables that correspond to these nodes. We show that,

$$
\mathbf{E}\left[X_{A} \cdot X_{B}\right] \leq \mathbf{E}\left[X_{A}\right] \cdot \mathbf{E}\left[X_{B}\right] .
$$

A simple induction then implies that $(X)_{v \in F_{t}}$ is negatively correlated: Pick any set $S \subset[1, n]$ and denote the indices in $S$ as $i_{1}, \ldots, i_{s}$. Given that the formula above is true for any disjoint subset, we have:

$$
\begin{aligned}
\mathbf{E}\left[\prod_{j=i_{1}, \ldots, i_{s}} X_{j}\right] & \leq \mathbf{E}\left[X_{i_{1}}\right] \cdot \mathbf{E}\left[\prod_{j=i_{2}, \ldots, i_{s}} X_{j}\right] \\
& \leq \mathbf{E}\left[X_{i_{1}}\right] \cdot \mathbf{E}\left[X_{i_{2}}\right] \cdot \mathbf{E}\left[\prod_{j=i_{3}, \ldots, i_{s}} X_{j}\right] \\
& \leq \\
& =\prod_{j=i_{1}, \ldots, i_{s}} \mathbf{E}\left[X_{j}\right] .
\end{aligned}
$$

Here, the inequalities hold by choosing $A=\left\{i_{1}\right\}$ and $B=S \backslash A$ in the first line.
Thus, we will now show that

$$
\mathbf{E}\left[X_{A} \cdot X_{B}\right] \leq \mathbf{E}\left[X_{A}\right] \cdot \mathbf{E}\left[X_{B}\right] .
$$

Observe that,

$$
\begin{aligned}
& \mathbf{E}\left[X_{A} X_{B}\right] \leq \mathbf{E}\left[X_{A}\right] \cdot \mathbf{E}\left[X_{B}\right] \\
\Longleftrightarrow & \mathbf{E}\left[X_{A} X_{B}\right]-\mathbf{E}\left[X_{A}\right] \cdot \mathbf{E}\left[X_{B}\right] \leq 0 \\
\Longleftrightarrow & \operatorname{Cov}\left(X_{A}, X_{B}\right) \leq 0 .
\end{aligned}
$$

Thus, we will show that $\operatorname{Cov}\left(X_{A}, X_{B}\right) \leq 0$.
If we condition on the set of tokens that each node receives, then $\left(X_{v}\right)_{v \in F_{t}}$ (and therefore $X_{A}$ and $X_{B}$ ) follows the hypergeometric distribution as we draw the received tokens without replacement. Thus, intuitively for any fixed distribution of tokens to fresh nodes, the conditioned distribution is therefore negatively correlated. To show this more formally, we need to to precisely condition $\left(X_{v}\right)_{v \in F_{t}}$ on the distribution of tokens. For a simpler presentations denote all tokens sent by $w$ as red tokens, all others are blue in the following. Let $R_{v} \in\{0,1\}$ the random variable that indicates that $v \in F_{t}$ received red token. Further let $B_{v}$ count the number of distinct blue tokens not received by $v$. Thus, the set $Y:=\left(R_{v} \cup B_{v}\right)_{v \in F_{t}}$ completely characterizes the distribution of tokens to nodes.

Using the closure properties of NA, one can show that the set $Y:=\left(R_{v} \cup B_{v}\right)_{v \in F_{t}}$ is NA: For each each red token $\theta$ we can define the binary variables $X(\theta, v)$. It holds that $X(\theta, v)=1$ if $\theta$ reaches $v$,
and 0 otherwise. Observe that for a fixed $\theta$ there is at most one $X(\theta, v)=1$ and all others are 0 as a token can only be received by one node. Thus, for each token $\theta$ the set $(X(\theta, v))_{v \in V_{t}}$ is NA. Further, if we fix a node $v$ the sum of all $X(\theta, v)$ is NA as well. The variable $R_{v}$ is now simply defined as

$$
R_{v}:= \begin{cases}1 & \text { if } \sum_{i=1}^{\tau} X\left(\theta_{i}^{w}, v\right)>0 \\ 0 & \text { else }\end{cases}
$$

Since it is monotonically increasing in $\sum_{i=1}^{\tau} X\left(\theta_{i}^{w}, v\right)$ the variable $R_{v}$ is also NA.
The same holds almost analogously for the blue tokens, however we need to adapt to the fact that we count the identifiers that were not received. Therefore, we define $\bar{X}(\theta, v):=1-X(\theta, v)$. In particular, we set $\bar{X}(\theta, v)=0$ if $\theta$ reaches $v$, and 1 otherwise. As $X(\theta, v)$ is NA, so is $\bar{X}(\theta, v)$ as it is monotone function. For each node $u \in V \backslash\{w\}$ now define $B_{v}^{u}$ to be binary variable indicating if any token of $u$ reached $v$. Thus, $B_{v}^{u}$ equals 1 if all $\tau$ tokens of $u$ missed $v$. We formalize this by setting:

$$
B_{v}^{u}:= \begin{cases}1 & \text { if } \sum_{i=1}^{\tau} \bar{X}\left(\theta^{u}, v\right)=\tau \\ 0 & \text { else }\end{cases}
$$

Again, $\sum_{i=1}^{\tau} \bar{X}\left(\theta_{i}^{u}, v\right)$ is NA and thus $B_{v}^{u}$ is increasing function of $\sum_{i=1}^{\tau} \bar{X}\left(\theta_{i}^{u}, v\right)$ it is NA, too. By summing over all $B_{v}^{u}$, we obtain the number of distinct tokens that did not reach $v$. Further, the variables $\left(B_{v}\right)_{v \in F_{t}}$ with $B_{v}:=\sum_{u \in V} B_{v}^{u}$ are NA as each variable is a monotone function on disjoint NA variables.

Now observe that all $X_{v}$ 's are independent of one another given $Y$. Further, the expected value $E\left[X_{v} \mid R_{v}, B_{v}\right]$ of each $X_{v}$ is dependent solely on the variables $R_{v}$ and $B_{v}$. In particular, each $E\left[X_{v} \mid R_{v}, B_{v}\right]$ monotonically rises in both $R_{v}$ and $B_{v}$ and is given by,

$$
\begin{aligned}
\mathbf{E}\left[X_{v} \mid Y=\left(R_{v}, B_{v}\right)\right] & =\operatorname{Pr}\left[X_{v} \mid Y=\left(R_{v}, B_{v}\right)\right] \\
& =\frac{R_{v}}{R_{v}+\left(m-B_{v}\right)}
\end{aligned}
$$

where $m=|V \backslash\{w\}|$ is the total number of nodes in round $t-(2 \lambda+5)$ that send blue tokens.
Obviously, receiving a red tokens raises the expectation of drawing a red token. However, it also holds that, the less blue token (i.e., more blue tokens that we don't receive) we receive the more likely it becomes to draw a red token. Thus, given two disjoint subsets $X_{A}, X_{B} \subset X$ we can view $E\left[X_{A} \mid Y\right]$ and $E\left[X_{B} \mid Y\right]$ as functions that monotonically rise in disjoint subsets $Y_{A}=\left(R_{A} \cup B_{A}\right)$ and $Y_{B}:=\left(R_{B} \cup B_{B}\right)$. With these observations, we can now show that $X$ is negatively correlated using only some technical arguments. By the law of total convariance, it holds:

$$
\operatorname{Cov}\left(\left(X_{A}, X_{B}\right):=\mathbf{E}\left[\operatorname{Cov}\left(X_{A}, X_{B}\right) \mid Y\right]+\operatorname{Cov}\left(\mathbf{E}\left[X_{A} \mid Y\right], \mathbf{E}\left[X_{B} \mid Y\right]\right)\right.
$$

We see that the first term is 0 since $\left(X_{v}\right)_{v \in F_{t}}$ is independent given $Y$. Note that the covariance of independent variables is always 0 by definition. Thus, it holds:

$$
\operatorname{Cov}\left(X_{A}, X_{B}\right)=\operatorname{Cov}\left(\mathbf{E}\left[X_{A} \mid Y\right], \mathbf{E}\left[X_{B} \mid Y\right]\right)
$$

It remains to show that this term is smaller than 0 . First, recall that $E\left[X_{A} \mid Y\right]$ and $E\left[X_{B} \mid Y\right]$ are monotonically increasing in $Y_{A}$ and $Y_{B}$. Thus, we can view them as monotonically increasing functions $f\left(Y_{A}\right):=\mathbf{E}\left[X_{A} \mid Y\right]$ and $g\left(Y_{B}\right):=\mathbf{E}\left[X_{B} \mid Y\right]$ for disjoint subsets of $Y$. Further, we showed $Y$ is NA and thus - by the closure properties of NA - its holds that $f\left(Y_{A}\right)$ and $f\left(Y_{B}\right)$ are NA, too. Therefore, the fact that $\operatorname{Cov}\left(f\left(Y_{A}\right), g\left(Y_{B}\right)\right) \leq 0$ follows from the definition of NA. And thus:

$$
\operatorname{Cov}\left(\mathbf{E}\left[X_{A} \mid Y\right], \mathbf{E}\left[X_{B} \mid Y\right]\right)=\operatorname{Cov}\left(f\left(Y_{A}\right), g\left(Y_{B}\right)\right) \leq 0 .
$$

This implies that $\left(X_{v}\right)_{v \in F_{t}}$ is negatively correlated.

Now, we need one last application of the Chernoff Bound. Using Lemma 31, we know that $E[X] \leq \delta$. Now observe that,

$$
\operatorname{Pr}[X>2 \delta]=\operatorname{Pr}\left[X>2 \frac{\delta}{E[X]} E[X]\right]
$$

Note that $\frac{\delta}{E[X]} \geq 1$ and we can therefore use the third bound from Lemma 4 for parameters bigger than 1. Thus, for $\delta \geq 3 k \lambda$ we get that:

$$
\begin{aligned}
\operatorname{Pr}[X \geq 2 \delta] & \leq \exp \left(-\frac{\delta}{3}\right) \\
& \leq \exp (-k \lambda)=\frac{1}{n^{k}}
\end{aligned}
$$

This shows the lemma as was claimed.

### 4.3 Congestion

Lemma 33 Algorithms Alg-LDS and Alg-Random have congestion of $O\left(\log ^{3} n\right)$ per node and round w.h.p.

Proof: We observe the number of messages due to Alg-LDS and Alg-Random by invocation of Alg-Routing. We observe the two subroutines Alg-LDS and Alg-Random separately.

1. In ALg-LDS each round every mature starts three routing requests for itself and three routing requests on behalf of each fresh node connected to it. Since there are at most $2 \delta$ fresh nodes connected to a mature node w.h.p., a given mature node starts $O(\log n)$ routing requests.
2. In Alg-Random each round every mature starts $O(\log n)$ tokens per round. Each token corresponds to one routing request.

Using Lemma 12 that each routing takes $O(\log n)$ rounds and Lemma 17 that for each routing request Alg-Routing has a congestion of $O(\log n)$ per round, Alg-LDS and Alg-Random together have congestion $O\left(\log ^{3} n\right)$ per round.

Now we observe the remaining operations performed each round.

1. Recall that each swarm is of size $O(\log n)$ w.h.p. Thus, during the introduction step in ALG-LDS each mature node introduces $O(\log n)$ nodes to their $O(\log n)$ neighbors. Resulting in a congestion of $O\left(\log ^{2} n\right)$ additional messages per node.
2. In Alg-Random each node, w.h.p., receives $O(\log n)$ tokens through the sampling algorithm and forwards them to fresh nodes. Additionally, each fresh node sends out $O(\log n)$ advertisements. Thus, altogether each node exchanges $O(\log n)$ messages.

Theorem 1 now follows from lemmas 20,24 and 33 ,

## 5 Future Work \& Conclusion

We presented an algorithm that maintains a structured overlay in presence of a $(2, O(\log n))$-late adversary. We permit $\alpha n$ deletions/additions over the course of $O(\log n)$ rounds. Note that this is exponentially higher than in [3] and [6. However, both their algorithms are not possible if the adversary has more recent knowledge of topology. This suggests a strong connection between an adversaries lateness with regard to the topology and permitted churn. For future work, one could consider finding an algorithm that tolerates a $(1, O(\log n))$-late adversary. Also one could consider a hybrid model where the adversary has an almost up-to-date information about some nodes but is more outdated with regard to others.

Further, we did not consider any kind of byzantine behavior. However, the approaches used by Fiat et. al. [9] could perhaps also be used with our overlay. Given, the overlay can handle byzantine nodes, further overlay problem, i.e. distributed agreement in the $(a, b)$-late setting, could also promising directions for future work.

## References

[1] John Augustine, Anisur Rahaman Molla, Ehab Morsy, Gopal Pandurangan, Peter Robinson, and Eli Upfal. Storage and search in dynamic peer-to-peer networks. In Proc. of SPAA, pages 53-62, 2013.
[2] John Augustine, Gopal Pandurangan, Peter Robinson, Scott T. Roche, and Eli Upfal. Enabling robust and efficient distributed computation in dynamic peer-to-peer networks. In Proc. of FOCS, pages 350-369, 2015.
[3] John Augustine and Sumathi Sivasubramaniam. Spartan: A framework for sparse robust addressable networks. In Proc. of IPDPS, pages 1060-1069, 2018.
[4] Baruch Awerbuch and Christian Scheideler. Towards scalable and robust overlay networks. In Proc. of IPTPS, 2007.
[5] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Proceedings of the 1st ACM Conference on Computer and Communications Security, pages 62-73, 1993.
[6] Maximilian Drees, Robert Gmyr, and Christian Scheideler. Churn- and dos-resistant overlay networks based on network reconfiguration. In Proc. of SPAA, pages 417-427, 2016.
[7] Devdatt Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. Random Structures \& Algorithms, 13(2):99-124, 1998.
[8] Michael Feldmann and Christian Scheideler. A self-stabilizing general de bruijn graph. In Proc. of SSS, pages 250-264, 2017.
[9] Amos Fiat, Jared Saia, and Maxwell Young. Making chord robust to byzantine attacks. In Proc. of $E S A$, pages 803-814, 2005.
[10] Robert Gmyr, Kristian Hinnenthal, Christian Scheideler, and Christian Sohler. Distributed monitoring of network properties: The power of hybrid networks. In Proc. of ICALP, 2017.
[11] Thorsten Götte, Kristian Hinnenthal, Christian Scheideler, and Julian Werthmann. Time-optimal construction of overlay networks. CoRR, abs/2009.03987, 2020.
[12] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. Ann. Statist., 11(1):286-295, 031983.
[13] Valerie King, Scott Lewis, Jared Saia, and Maxwell Young. Choosing a random peer in chord. Algorithmica, 49(2):147-169, 2007.
[14] Valerie King and Jared Saia. Choosing a random peer. In Proc. of PODC, pages 125-130, 2004.
[15] Michael Mitzenmacher and Eli Upfal. Probability and Computing. Cambridge University Press, 2005.
[16] Moni Naor and Udi Wieder. Novel architectures for P2P applications: the continuous-discrete approach. In Proc. of SPAA, pages 50-59, 2003.
[17] Andréa W. Richa, Christian Scheideler, and Phillip Stevens. Self-stabilizing de bruijn networks. In Proc. of SSS, pages 416-430, 2011.
[18] Christian Scheideler. Probabilistic Methods for Coordination Problems. PhD thesis, Paderborn University, 2000.
[19] Christian Scheideler. How to spread adversarial nodes?: rotate! In Proc. of STOC, 2005.
[20] Daniel Stutzbach and Reza Rejaie. Understanding churn in peer-to-peer networks. In Proc. of SIGCOMM, pages 189-202, 2006.
[21] David Wajc. Negative association-definition, properties, and applications. 2017.


[^0]:    *A preliminary version appeared in the proceedings of the 33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2019.
    ${ }^{\dagger}$ Department of Computer Science, Paderborn University, Germany, E-mail:thgoette@mail.upb.de. This work is partially supported by the German research council (DFG) in the context of the collaborative research center "On-the-Fly Computing" (SFB 901)
    ${ }^{\ddagger}$ Chair of Management Science, RWTH Aachen, Germany. E-mail:vipin.rv@oms.rwth-aachen.de. This work is partially supported by the German research council (DFG) Research Training Group 2236 UnRAVeL.
    ${ }^{\S}$ Department of Computer Science, Paderborn University, Germany, E-mail:scheidel@mail.upb.de

[^1]:    ${ }^{1}$ Throughout this paper w.h.p. means with probability $\left(1-\frac{1}{n^{k}}\right)$, where $n$ is the number of nodes and $k$ is a tunable constant
    ${ }^{2}$ Synchronicity is a standard assumpution in the related work as nodes need to react to the adversary's changes in a timely manner.

[^2]:    ${ }^{3}$ In principle our algorithm in Section 4 could be extended to tolerate $O(p o l y \log n)$ joins per node as in [3]. We chose a constant because a higher number of joins would only increase the number of messages by a poly-logarithmic factor and does not introduce further algorithmic challenges.

[^3]:    ${ }^{4}$ Note that we do not seek to optimize message complexity

[^4]:    ${ }^{5}$ For convenience we assume throughout this work that $\lambda$ is an integer

[^5]:    ${ }^{6}$ Note that since $c$ is a constant and $\lambda \in O(\log n)$, therefore for big enough $n, \frac{c \lambda}{n} \leq 1$.

[^6]:    ${ }^{7}$ Given that $\alpha n$ nodes may fail in single round and we want to route each message on the first try, it reasonable that one needs $\mathcal{O}(\log n)$ copies of a message each round to ensure the survival of at least one w.h.p.

