Stability of P2P Networks Under Greedy Peering

Lucianna Kiffer[1] and Rajmohan Rajaraman[2]

[1] ETH Zurich lkiffer@ethz.ch
[2] Northeastern University r.rajaraman@northeastern.edu

Abstract. Historically, major cryptocurrency networks have relied on random peering choice rules for making connections in their peer-to-peer networks. Generally, these choices have good properties, particularly for open, permissionless networks. Random peering choices however do not take into account that some actors may choose to optimize who they connect to such that they are quicker to hear about information being propagated in the network. In this paper, we explore the dynamics of such greedy strategies. We study a model in which nodes select peers with the objective of minimizing their average distance to a designated subset of nodes in the network, and consider the impact of several factors including the peer selection process, degree constraints, and the size of the designated subset. The latter is particularly interesting in the context of blockchain networks as generally only a subset of nodes are miners/content producers (i.e., the propagation source for content), or are running specialized hardware that would make them higher performing.

We first analyze an idealized version of the game where each node has full knowledge of the current network and aims to select the d best connections, and prove the existence of equilibria under various model assumptions. Since in reality nodes only have local knowledge based on their peers' behavior, we also study a greedy protocol which runs in rounds, with each node replacing its worst-performing edge with a new random edge. We exactly characterize stability properties of networks that evolve with this peering rule and derive regimes (based on number of nodes and degree) where stability is possible and even inevitable. We also run extensive simulations with this peering rule examining both how the network evolves and how different network parameters affect the stability properties of the network. Our findings generally show that the only stable networks that arise from greedy peering choices are low-diameter and result in disparate performance for nodes in the network.

1 Introduction

The use of cryptocurrency networks has exploded in recent years, with billions of dollars being transferred daily across various platforms [6]. As these networks grow, it becomes increasingly important to understand the factors that impact their structure and stability. Traditionally, networks like Bitcoin and Ethereum[3] have implemented peer selection strategies which aim at randomizing peer choices to avoid attacks [14,24,13], and favor properties such as low diameter [5,10] for fast information propagation. Recently in Ethereum's new Proof-of-Stake network, the Beacon network, some clients have taken a new approach of greedily choosing peers based on performance [21] where nodes drop some fraction of worst-behaving peers. In this work, we analyze a simplified model of such greedy peering strategies.

This paper explores the dynamics of greedy peering strategies, specifically the case where nodes aim to optimize their distance to all nodes in the network or to a special subset of high-value nodes (which we term the *miners*). We are interested in characterizing how heterogeneous peering preferences affect the evolution and stability of the network. Our study is driven by the following questions:

1. Do greedy peering strategy games have stable networks and can we characterize the conditions for stability?
2. What are the properties of stable networks? When stability is reachable, how do networks converge to a stable state?
3. Do protocol parameters (e.g., degree constraints, number of miners) alleviate or exacerbate stability dynamics?

1.1 Our results

We consider networks of n nodes where $m \leq n$ nodes are special and are labeled as *miners*; these special nodes represent miners in a cryptocurrency network or content generators in a gossip-based dissemination

[3] Historically the two largest by market cap [6].

network. The primary objective for each node is to be "close" to the miners in the network; we assign a score for each node as the average hop-distance to the miners in the *undirected network* formed by all the edges. The P2P network topologies that evolve in the process depend on a number of factors, including the fraction of the network that are miners, and nodes' out- and in-degree constraints.

We begin by evaluating an idealized version of a non-cooperative game where nodes choose their optimal peers to minimize their own score (i.e., average distance to the miners).

- **Existence of stable networks:** We show that *there always exist pure Nash equilibria (or stable networks) when nodes have unbounded in-degree*. We also prove that *Nash equilibria always exist for a non-trivial infinite family of networks with bounded in-degree*; the general case with bounded in-degrees is open. These results are presented in Section 4.1.

While the idealized game provides a sense of stable network topologies that can arise, the game is hard to implement in practice since nodes do not have global knowledge of the network and cannot make arbitrary peering choices.

Our main results concern a natural peering protocol, which we call *Single-Exploratory-Greedy*, where nodes have knowledge of only their neighbors' scores and can add an exploratory random link to replace an existing link. For this protocol, we consider a notion of stability in the network with respect to stable edges (edges that will never be dropped).

- **Conditions for stability:** We find that under the Single-Exploratory-Greedy protocol, *often the only stable networks involve a centralized, well-connected core of nodes (the miners), with all other nodes directly connecting to this core*. We also show that *when nodes are faced with ties on which connection to drop, the Tie Breaker Rule is a determining factor on whether any stable topologies exist*; for instance, we show the impossibility of stability with two such rules. Our detailed analysis of Single-Exploratory-Greedy is in Section 4.2 with a summary of results in Table 1.
- **Reaching stability:** We also show that when the number of miners is capped (to at most the node out-degree capacity), if a stable network exists, networks starting from any initial state will *always eventually* reach a stable state.

The bulk of our theoretical results are for networks with unbounded in-degree. To explore how and if bounded networks differ in their evolution to stable topologies, we run simulations of the Single-Exploratory-Greedy protocol in diverse scenarios (see Section 5).

- **Connectivity:** Our simulations largely support that in networks where a subset of the nodes are miners, even with bounded in-degree, the miners become more connected to one another and more centralized in the network over time (i.e., lower miner diameter and eccentricity), and that this behavior is quicker in networks where miners have a higher probability of connecting (e.g., smaller networks, larger inbound capacity, more than one exploratory edge).
- **Fairness:** Additionally, across the board, our simulations show that though Single-Exploratory-Greedy lowers network properties like average distance to miners, diameter, and average eccentricity, this often comes at the price of fairness, where some nodes are worse off in the network (e.g., have a higher average distance to miners than in the random network) and that the network-level average advantage is likely due to a few nodes being much better off and skewing the average.

2 Related Work

We aim to understand the topological impact of a greedy peering protocol on a peer-to-peer network with a subset of preferential nodes everyone is trying to get close to. Most major existing P2P networks of this kind, namely cryptocurrencies, implement peering protocols that approximate low-diameter random networks [5,10]. Topological analysis of these networks, however, is hard as they tend to obfuscate peer connections to avoid attacks [14,24,13], and existing measurement techniques are prohibitively expensive or invasive [8,20], thus they tend to be run on test networks as a proof of concept. In [8], Delgado et al. map the topology of the Bitcoin Ropsten test network and compared it with simulated random networks of the same size and similar degrees. They find the Ropsten network to have some behavior similar to networks in our simulations (e.g., a small number of concentrated central nodes).

An inspiration for the Single-Exploratory-Greedy of this paper is the Perigee protocol of Mao et al. [23], which also implements a greedy peering protocol. The authors consider a network of all miners in an Euclidean space and show theoretically that nodes choosing neighbors geographically close to them results in network distance that is close to the optimum distance (shortest path is a constant away from the Euclidean distance). They show via simulations that the Perigee protocol approximates Euclidean distances between all nodes. In their simulations, they show that as processing time at each node increases, the performance of Perigee and the random network converge. In our work, this is the regime we are interested in studying: negligible link latency in comparison to node processing latency such that the properties of the overlay topology (and hop distance) are more important than the underlying communication link latency. We also focus on networks where a subset of the nodes are miners, as this is often the case [26,18]. Node processing times and congestion have also been a recent focus in security analyses of blockchain consensus protocols [27,16].

There is a growing body of work on peering protocols in cryptocurrency networks. Recent work has shown that variations of the Perigee protocol run by a single entity can work well as a strategy for reducing a node's own latency to the source of transactions in a network [33,3]. In [30], Park et al. use their own measurements of Bitcoin to propose an IP-based distance peer selection protocol to reduce round trip time of peers. Though this is in the same spirit of Perigee-like greedy protocols, using IP location is notoriously unreliable [31,12], and tricky as a basis for routing as the Internet does not satisfy the triangle inequality [22]. Other proposals for overlay networks include randomly replacing edges to maintain approximate sparse random graphs [36], and more structured proposals such as one based on the DHT Kademlia [32], and a hypercubic overlay network [2], as well as others [34].

Our game-theoretic formulation falls under the category of network creation games that have been extensively studied by researchers in computer science, game theory, and economics. One of the earliest works in this space is due to Jackson and Wolinsky who introduced a model in which there is an underlying cost for each link, and studied tradeoffs between stability and efficiency [15]. Demaine et al studied the price of anarchy in network creation, quantifying the overhead of the worst stable network, when compared with the social optimum [9]. Many different models based on cost, direction of links, degree constraints, and initial conditions have been introduced, with the aim of capturing diverse scenarios including social networks and the Internet [4,7,19,25]. Our game-theoretic model is different from models studied in past work in the following way. Previous work has considered either directed networks with the node actions being given by directed edges (e.g., [19]) or undirected networks with the node actions being given by undirected edges (e.g., [25]). In contrast, P2P networks underlying blockchains are better modeled by processes in which nodes select directed out-going edges to peers under some capacity constraints, while the latency incurred in the resulting network is independent of the edge direction. Previous models in network creation games do not consider this important distinction on how links are created and their impact on delays. As a result, the stability and convergence properties of past models differ from the results we derive for our model. Furthermore, past work on network creation games is primarily on the nature of Nash equilibria; while we also study equilibria for our model, our focus is on the analysis of a greedy peering process that captures the peering approaches of real P2P networks.

3 Model

We consider a network with n nodes, the first m of which are miners, with each node having d out-going edges and at most $d_{in} \leq n$ incoming edges. Let G denote the graph consisting of the nodes and the directed edges. We use (i, j) to denote the directed edge from i to j. For convenience, we also use $e_{i,j}$ as the indicator variable for whether either of edges (i, j) or (j, i) exist; so, $e_{i,j}$ is 0 (resp., 1) if neither (i, j) nor (j, i) exist (resp., if (i, j) exists or (j, i) exists) in G. Thus, $e_{i,j} = e_{j,i}$. We use the terms peer and neighbor interchangeably, and edge and connection interchangeably. We also refer to d_{in} as the inbound connection capacity.

The **Single-Exploratory-Greedy** protocol operation proceeds in rounds, each of which consists of the following steps:

1. In an arbitrary order, each node i selects at random a peer j s.t. currently $e_{i,j} = 0$ (thus preventing bi-directional edges) and j has not reached it's inbound capacity, and adds (i, j) to G (thus setting $e_{i,j} = 1$). If i has an edge to or from every node, it adds no new edges.

2. Every node i updates its score $score_G(i)$ as follows. If $dist_G(i,j)$ denotes the shortest distance between i and j in the *undirected* version of G, then

$$score_G(i) = \sum_{j \in miners} \frac{dist_G(i,j)}{m},$$

which is the average distance to all miners. A miner's distance to themselves is 0, thus the minimum score for a miner is $\frac{m-1}{m}$ and the minimum score of a non-miner is 1.
3. Each node i with d out-going edges after the above steps, drops an out-going edge (i,j) (i.e., sets $e_{i,j} = 0$), where j is the out-neighbor of i with the maximum score; if i does not have d out-going edges then it does not drop a peer. If two or more peers have the highest score, then a Tie Breaker Rule as defined below is used to pick one peer.

Definition 1 (Tie Breaker Rule). *We define the Tie Breaker Rule as the rule used for deciding ties when a node has multiple peers sharing the highest score in step (3). We consider nodes deciding ties according to the following possible rules: uniformly at random (Random), First-In-First-Out (FIFO), Last-In-First-Out (LIFO), or a global ordering where all nodes are ranked in some order that does not change during the protocol and the node with highest rank wins (Global-Ordering).*

Definition 2 (Stability). *We define an edge as* ***stable*** *at some round r, if its value at the end of the round (after step 3), does not change for all future rounds. Formally, let $e^r_{i,j}$ be the value of $e_{i,j}$ at the end of round r. Edge $e_{i,j}$ is stable at round r, if $\forall r' \geq r$, $e^{r'}_{i,j} = e^r_{i,j}$.*

A ***network-stable topology*** *is one where all nodes have $d-1$ stable outgoing edges. Formally, G is network-stable at round r if $\forall i,j \leq n, r' \geq r$ $e^{r'}_{i,j} = e^r_{i,j}$. A network-stable topology is therefore an equilibrium since each node has one outgoing edge they keep replacing with a random node, but the choice of this edge will never cause a stable edge to be dropped.*

A ***miner-stable topology*** *is one where all edges between miners are stable, and therefore no further edge between miners will arise outside of the d-th edge. Formally, graph is miner-stable at round r if $\forall i,j \leq m, r' \geq r$ $e^{r'}_{i,j} = e^r_{i,j}$.*

Note that for a topology to be network-stable or miner-stable, we require all nodes to have at least $d-1$ outgoing edges since a node will never drop an edge if they have less than d. Additionally, the following event arises often in our arguments related to stability.

Definition 3 (Special Event). *We define Minimum-Score-Event$_i$ as the event where at step (1) of some round, all miners who were not yet connected to some node i, become connected to i. At step (2) of the protocol, node i thus has its minimal score ($\frac{m-1}{m}$ for miners, and 1 for non-miners). Assuming nodes have no inbound connection capacity, there is always some probability this event happens for a given i in the next round.*

This event goes to show that it is always possible for any node to come to have its minimal score in any round, thus, most of our proofs involve showing that stable connections must be to minimal scoring peers so that they are not replaced.

Model motivation. In our model (and in practice), nodes do not know if a peer is a miner or not, all they know is a peer's relative performance. Thus, both miners and non-miners are trying to minimize their distance to all miners via the score function. While for simplicity, in our model, we use hop distance to score peers, in practice, network propagation delays can be used locally to score peers against each other. We also adopt a graph model where the set of connections are made in a directed fashion, while their performance is evaluated over the undirected edges. This mirrors the behavior of real P2P networks, where nodes generally have a set inbound and outbound connection cap, and will accept connections as long as they have slots (allowing new nodes who join the system to find peers) but will be more selective with their out-going edges (e.g., to prevent DoS). When it comes to information propagation however (e.g., gossiping of blocks and transactions), all peers are treated the same and thus message delay is dependent only on undirected hops. Additionally, we omit a peer discovery process and instead assume nodes can sample connections at random from all nodes in the network[4]. We also consider a greedy peering process that proceeds in rounds, though within each step there is no coordination between nodes (except the assumption that in step 1 no two nodes choose each other). These simplifying assumptions allow for an initial clean theoretical evaluation.

4 Theoretical Analyses

Our goal is to analyze Single-Exploratory-Greedy and the underlying network formation process. To begin, we consider an idealized game-theoretic version of Single-Exploratory-Greedy where nodes have full knowledge of the graph and can choose all peers to optimize their score, and explore the equilibria of this game. We then analyze the Single-Exploratory-Greedy protocol, studying the stability of the networks that arise during the protocol.

4.1 An idealized game-theoretic model

We consider first an idealized game-theoretic model in which each node wants to select their outgoing peers so as to minimize the average distance to a miner. A natural approach is via the concept of Nash equilibrium [29]. Let d denote the out-degree of each node. For a given node v, an action is a set of at most d nodes of the network representing the out-neighbors that v selects. Note that in this idealized game, nodes can choose *all* new outgoing edges. Recall that while we distinguish between in- and out-edges, the distances among nodes and the average diameter are based on the graph that is obtained by viewing each edge as undirected.

Let $\mathcal{G}(n, m, d)$ (resp., $\mathcal{G}(n, m, d, d_{in})$) denote the *uncapped* (resp., *capped*) *game* with n nodes, $m \leq n$ miners, out-degree d, and unbounded in-degree (resp., in-degree at most d_{in}). We first show that the uncapped game always has a pure Nash equilibrium, a stable topology in which no node wants to change its out-neighbors.

Lemma 1 (Every Uncapped Game has a Stable Network). *For all n, m, d, $\mathcal{G}(n, m, d)$ has a Nash equilibrium.*

Proof. Consider any graph G with the following properties: (i) there exists a node r such that every node in G has an edge to r; (ii) every edge is directed to a miner; (iii) there does not exist any pair of nodes u and v such that (u, v) and (v, u) are both present in G; and (iv) if a node v has less than d outgoing edges, then it has an edge to or from every miner.

We claim that G is a Nash equilibrium. We compute the score for each node in G. The score of miner r equals $(m-1)/m$ since every other miner is at distance 1; this is the lowest score possible and no action of r can improve it. Let $u \neq r$ denote any other miner, and let i_u denote the number of miners that have edges to u. Since r is adjacent to every other node, u is at distance at most 2 from any other miner. Given the choices of all other nodes, a best-response action of u is to set all its out-edges to miners that do not have an edge to u; the best score achievable equals $\max\{(d + i_u) + 2(m - 1 - i_u - d), (m - 1)\}/m$. This is precisely what is achieved in G given the conditions: by (ii), all edges are directed to a miner; by (iii) u does not put any out-edges to a miner that already has an edge to u; finally, by (iv), if u does not use all of its d edges, then it already has a score of $(m - 1)/m$.

We finally consider a non-miner v. By the conditions, all of the edges of a non-miner are to miners. Its distance to $\min\{d, m\}$ miners is 1 and its distance to every other miner is 2, leading to an optimal possible score of $\max\{(d + 2(m - d))/m, 1\}$. Thus, v has no incentive to change its out-edges. We have thus shown that G is a pure Nash equilibrium. □

We next consider the network formation game with bounded in-degree. Since the endpoints of the out-edges are constrained by this bound, reasoning about stable networks is much more challenging. We first show that the simplest network formation games in the capped setting, the family $\mathcal{G}(n, m, 1, d_{in})$, have pure Nash equilibrium for all d_{in} and infinitely many choices of n and m. We sketch the graphs constructed in the following proofs in Figure 1.

[4] In reality, information of new nodes in the network needs to propagate and, though this generally happens quickly [11], it is possible that propagation of information on new nodes impacts how nodes cluster in the network.

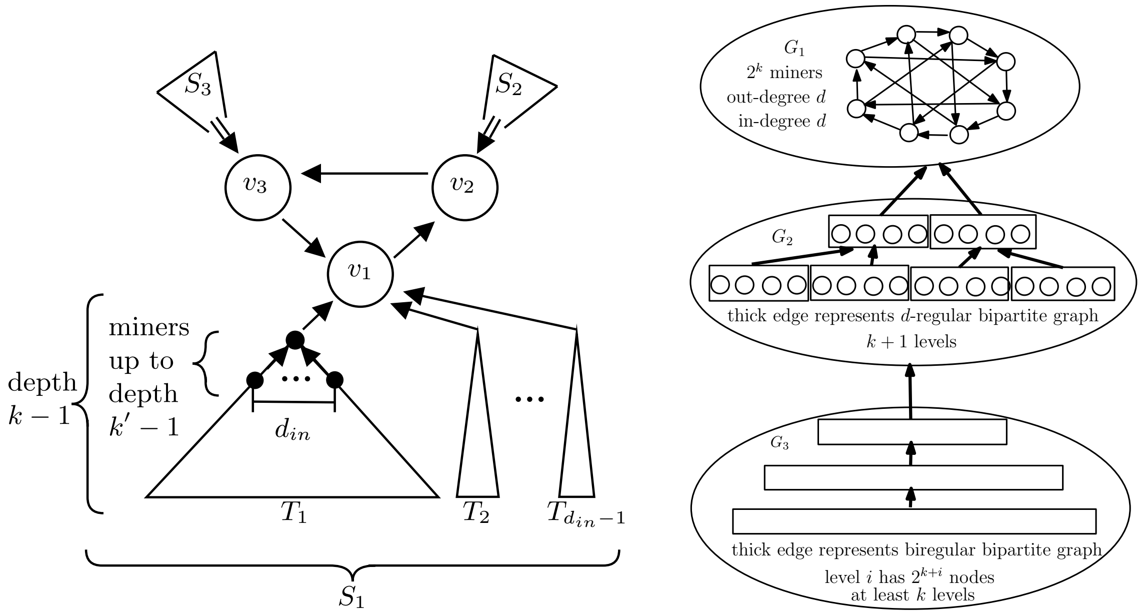

(a) Sketch from Lemma 2 (b) Sketch from Lemma 3

Fig. 1: For the existence proofs of Lemmas 2 and 3 we construct the above network families. Both networks are essentially layered, and the crux of the proof is that the miners are positioned in the "top layer(s)" with the non-miners occupying the remaining layers, such that all nodes except the non-miners in the last layer have full in-capacities. The nodes in the last layer have the highest score and no node higher up in the layers would give up their position to connect to a node in the bottom layer.

Lemma 2 (Stable Networks in Capped Games with Unit Out-Degree). *For any* $n = 3 \cdot d_{in}^k$ *and* $m = 3 \cdot d_{in}^{k'}$ *for any integers* k*,* k'*,* $k \geq k' \geq 0$*,* $\mathcal{G}(n, m, 1, d_{in})$ *has a pure Nash equilibrium.*

Proof. Note that n and m are at least 3. Let C be a directed triangle over three miners v_1, v_2, and v_3. Partition the remaining nodes into three equal-sized sets S_1, S_2, and S_3. By the assumption on n, each S_i has $d_{in}^k - 1$ nodes and $d_{in}^{k'} - 1$ miners. Split each S_i into $d_{in} - 1$ balanced complete d_{in}-ary directed trees, with every edge directed to the parent. Note that a balanced complete d_{in}-ary tree of depth $h-1$ has $\frac{d_{in}^h - 1}{d_{in} - 1}$ nodes. Furthermore, assume that all nodes in each d_{in}-ary tree T_j up to depth $k'-1$ are miners (there are exactly $\frac{d_{in}^{k'} - 1}{d_{in} - 1}$ such nodes in each T_j). Let $r_{i,j}$ for $1 \leq j \leq d_{in} - 1$ be the root of each tree T_j in S_i. Consider the network G consisting of the union of C, S_1, S_2, S_3, and the edges $\{(r_{i,j}, v_i) : 1 \leq i \leq 3\}$. We show that G is a pure Nash equilibrium.

For each v_i, consider the directed tree Tv_i made up of the trees in S_i and the edges from their roots to v_i. By symmetry, the scores of all nodes at a given depth in Tv_i are the same. Let δ_1 and δ_2 be integers such that $k \geq \delta_1 > \delta_2$. Consider two nodes n_1 and n_2 in S_i at depths δ_1 and δ_2, respectively. The score of n_1 is strictly greater than that of n_2. This is because a node v at some depth δ has a lower score than any of the nodes in the sub-tree below it since there are twice as many miners not in Tv_i as are in Tv_i, and v is in the only path of its sub-tree to those miners. This also implies that for any leaf ℓ and any non-leaf node v, the score of ℓ is at least the score of v. Given the symmetry of the leaves, this applies for all n_1, n_2 at depths δ_1 and δ_2 in different Tv_i trees.

To complete the proof, we consider each node and verify that their out-edge in G gives the node the best score, given the choices of other nodes. By construction, every non-leaf node in G has a full in-degree of d. Hence, the action of any node v in the game is to either have their current out-edge in G or to direct the out-edge to one of the leaf nodes. If v selects a leaf node ℓ, note that by symmetry this node has score s at least that of v. If v is not a miner, then its distance to all miners will be one more than that of ℓ, so the new score of v is $s + 1$.

If v is a miner at depth h, we consider the change in its score as new_score(v)−old_score(v) and prove this change is positive. Note first that v's distance to itself and all miners in the sub-trees that point to v remains the same, so a difference of 0. WLOG, assume v was in Tv_1 and chose to connect to a leaf in Tv_2. Thus v's distance to miners in Tv_3 has only increased by $k+1-h$. Next, let v' be the new ancestor of v in Tv_2 at height h. Since v and v' were at the same height in G, by the symmetric property of G, they had the same cumulative distance to miners in Tv_1 and Tv_2 who are not in the sub-trees of v or v'. Thus v's score to these miners has increased by its distance to v' which is again $k+1-h$.

The only remaining miners to consider in the change in score is from v to the miners in the sub-tree of v' which may now be closer (i.e., take away from the change in score). In the original structure, v was a maximum $2h+1$ from v' ($2h$ if v and v' are in the same tree), thus these miners are now at most $k+h+1$ from v. If we consider only the sub-trees of G whose roots are at depth h, let X be the number of miners in each of these sub-trees. There are $(d-1)*d^{h-1}$ such sub-trees in each S_i. Then new_score(v)−old_score(v) is at least the score difference for these sub-trees (since we are just under-counting miners who are now further away from v), we get that

$$\text{new\_score}(v) - \text{old\_score}(v) \geq (k+1-h)\cdot\left[3\cdot(d-1)\cdot d^{h-1}-2\right]\cdot X - (k+h+1)\cdot X.$$

Let's assume this score difference is non-positive. Remembering that $d \geq 2$, we can simplify the above inequality to be

$$2^{h-1} \leq (d-1)d^{h-1} \leq \frac{3k+h-3}{3k-3h+3}$$

which is not satisfiable. Thus the score difference is strictly positive, meaning v would not choose to change its out-edge. □

We next extend Lemma 2 to capped games with arbitrary out-degrees, establishing the existence of stable networks in much more general settings in Lemma 3. At a very high level, the proof establishes stability of networks with a certain structure: a strongly connected core G_1 consisting of the miners, a hierarchical layered network G_2 with nodes in its top layer having directed edges to G_1, and a more loosely structured layered network G_3 with nodes in its top layer having directed edges to G_2 and nodes in its bottom layer having no incoming edges. The nodes in the bottom layer of G_3 are farthest from the miners while the in-degree of every other node in the network equals the in-capacity, which ensures that the network is stable. Figure 1(b) illustrates this proof.

Lemma 3 (Stable Networks in Capped Games with Arbitrary Out-Degree). *For integers $k \geq 0$, $d \geq 2$, $m \leq 2^k$, and $n \geq 2^{3k}$, $\mathcal{G}(n, m, d, 2d)$ has a pure Nash equilibrium.*

Proof. We consider a network G that has three components. The first component G_1 is a graph with its vertex set being all the miners, each node with an out-degree of d to other miners, leaving $d2^k$ incoming capacity.

The second component G_2 consists of $k+1$ levels of nodes, with level $i \geq 0$ consisting of set L_i of 2^{k+i} nodes. For each i, we partition the 2^{k+i} nodes into 2^i groups of 2^k nodes each; we label these groups $S_{i,0}, \ldots, S_{i,2^{i+1}-1}$. For any non-negative integer $j_{k-1} < 2^k$, the node groups S_{i,j_i}, $j_i = \lfloor j_{i+1}/2 \rfloor$, $-1 \leq i < k$, include a butterfly network. That is, the edges between S_{i,j_i} and $S_{i+1,j_{i+1}}$, with $j_i = \lfloor j_{i+1}/2 \rfloor$ include the ith level of an 2^k-input butterfly network. This can be set up since $d \geq 2$. This ensures that every node in L_k has a path of length k to every node in $S_{1,0}$. (Instead of using a butterfly network, we could also have random biregular bipartite graphs between consecutive levels, achieving the same distance property between L_k and $S_{1,0}$ with high probability.) We can then place all remaining edges from each L_i arbitrarily to nodes in L_{i-1}. We then have all edges from $S_{1,0}$ go into G_1, thus all nodes at level L_k have the same score as all their paths to the miners converge at L_0.

The final component G_3 consists of at least k levels of nodes, which we number from k. For level $i \geq k$, the set L_i consists of 2^{k+i} nodes, with the last level ℓ possibly having fewer than $m2^\ell$ nodes. There is an arbitrary bipartite graph connecting L_i to L_{i-1} with d outgoing edges for each node in level i and $2d$ incoming edges for each node L_{i-1}.

We now argue that G is a stable network. We make three observations about distances, which are critical in establishing that no node benefits by changing any of its outgoing edges.

First, for any $0 \leq i < k$, for any node u in level i and miner v, the distance between u and v is at most $2k - i + \delta_v$, where δ_v is the minimum, over all nodes in L_0, of the distance between the node and v. This follows from the fact that any node in L_k has a shortest path to every node in L_0; consequently, any node in G_2 has a path to v by concatenating a subpath of length $k - i$ to the closest node in L_k, followed by a subpath of length k to the node in L_0 that is closest to v, followed by a subpath of length δ_v to v.

Second, for any $i \geq k$, for any node u in level i and miner v, the distance between u and v is exactly equal to $i + \delta_v$. This holds since any node in L_i has a path to v by concatenating a subpath of length $i - k$ to an arbitrary node in L_k, followed by a subpath of length k to the node in L_0 that is closest to v followed by a subpath of length δ_v to v.

Finally, we note that for any leaf node u and any miner v, the distance between u and v is at least $\ell + \delta_v$. To complete the proof, we observe that for any u in G and any miner v, the distance between u and v is at most $\max\{i, 2k - i\} + \delta_v \leq \ell + \delta_v$; consequently, replacing any out-edge with an out-edge to a leaf will not decrease its cost. Therefore, G is a stable network. □

Lemma 3 can be further extended to allow for more general relationships between the in-degree and the out-degree (e.g., when $d_{in} = cd$ for any integer constant $c \geq 3$). We do not know, however, if stable networks exist for all choices of n, m, d, and d_{in}. We also note that the proofs of Lemmas 1 through 3 are all existence proofs in the sense that they present *specific* families of stable networks. There are, however, many other families of stable networks, as suggested by the proofs. In particular, the proof of Lemma 3 presents a layered structure for stable networks that accommodates a variety of different topologies.

4.2 Single-Exploratory-Greedy

Our results in Section 4.1 indicate that there exist Nash equilibria in the full-information setting, where every node knows the current network topology at all times. In practice, however, nodes operate with information only about their own peers. We consider now the Single-Exploratory-Greedy protocol defined in Section 3 and analyze networks that evolve in this process. Instead of equilibria, we consider miner and network stability as defined in Section 3. We focus on two questions: does there always exist a stable topology? what properties do stable topologies have? Table 1 summarizes our results and the impact of the Tie Breaker Rule on the ability of the network to stabilize. We begin by defining conditions for miner stability.

	$m < 2d$		$m \geq 2d$	
	miner-stable	network-stable	miner-stable	network-stable
Global-Ordering/LIFO	✓ Prop. 1	✓ Lem. 5	✓ Lem. 4	✗ Lem. 6
Random/FIFO	✓ Prop. 1	✗ Lem. 7	✗ Prop. 2	✗ Lem. 7

Table 1: Summary of our results of Section 4.2 for the existence and impossibility of miner-stable and network-stable topologies in the unbounded Single-Exploratory-Greedy game. Note for network-stability, we assume $n \gg d$, as discussed in their proofs.

Proposition 1 (Clique Miner Stability for Small m). *If $m < 2d$ there exists a miner-stable topology for any Tie Breaker Rule.*

Proof. Since $m < 2d$, there are enough edges between miners such that they can connect in a clique with each miner having at most $d - 1$ outgoing edges to any other miner. Each miner i has a distance 0 to itself and a distance 1 to each of the other miners, thus their weighted distance to miners in the clique is $w_i = \frac{m-1}{m}$. Any node i drops a connection to the peer with the largest score among the peers of i. Any non-miner node h must have a distance of at least 1 to each miner, thus $score(h) \geq 1 > \frac{m-1}{m}$, thus a non-miner score will not tie with a miner score, therefore miners will only ever drop a peer who is not a miner. Since all miners are already connected, a miner will never establish a new connection to another miner. □

Lemma 4 (Global-Ordering and LIFO Miner Stability). *For all m and d, there exists a miner-stable topology for the cases where the Tie Breaker Rule is Global-Ordering or LIFO.*

Proof. For $m < 2d$, by Proposition 1 a clique is miner-stable for any tie-breaking rule. For all other values of m and d, we first examine the Global-Ordering rule. Suppose the first $d-1$ miners (by global ordering) connect in a clique and all other miners have their $d-1$ edges to the first $d-1$ miners. The first $d-1$ miners are therefore connected to all miners and have a minimum score $\frac{m-1}{m}$. In a round, these $d-1$ miners will not drop one another since they have the lowest score and priority ranking, and all other miners will not drop these first $d-1$ miners for the same reason. Since the first $d-1$ miners are connected to all miners, they will never develop a new connection to a miner. In a round, two miners i, j could connect to each other via their d-th outgoing edge. Since either miner could at best have the same score, but worst ranking as the first $d-1$ miners, this new edge will not replace the existing $d-1$ out-going edges to the first $d-1$ miners.

Next, consider the LIFO Tie Breaker Rule. Consider the graph where some set of $d-1$ miners connect in a clique, and all other miners have outgoing edges to these miners. Let the non-clique miners have their d-th edge to some non-miner. Note that the clique does not have any outgoing edges to the non-clique miners and the non-clique miners have no edges among themselves. Since all edges between miners are going into the clique, for any new edge to replace one of these edges, it would need to have an equal score but would be the newest edge, by LIFO it would be dropped. The clique is connected to all miners so it can't make a connection to any other miner. Miners not in the clique thus have $d-1$ stable edges. □

We note that the networks presented in the proof of Lemma 4 are not stable under FIFO (or Random) as there is always a chance all miners connect to some non-clique miner and this miner would always (or with some probability) replace an existing miner in the clique. We extend this logic to the following claim.

Proposition 2 (Impossibility of Miner-Stability under FIFO or Random for Large m). *If $2d \leq m$, and ties are broken by FIFO or randomly, there is no miner-stable topology.*

Proof. Assume there is some miner-stable topology. This means all edges between miners are either stable or will always be dropped at the end of the round (i.e., are the maximum-score edge after the Tie Breaker Rule). Since $2d \leq m$, miners cannot form a *stable* clique, thus there is some miner m_i that is not connected to all miners. There is some probability that in the next round, all miners connect to this miner with at least one miner m_j having a new out-going edge to m_i. Miner m_i now has score $\frac{m-1}{m}$; so, m_j will not drop the edge to m_i since it is either smaller than some existing edge or ties with all edges and by FIFO (or random) Tie Breaker Rule replaces some edge (has a probability of replacing some edge). □

We now derive conditions under which stable networks exist and conditions where stability is impossible. For the following cases, we generally assume $n \gg d$, and $n > 2d^2$ for proofs of instability[5]. Recall that since all nodes have at least $d-1$ outgoing edges at the start of the protocol, nodes always have at least $d-1$ outgoing edges during any portion of the protocol.

Lemma 5 (Stable Topology for Small m). *For all $m < 2d$ with Tie Breaker Rules LIFO or global ordering, there exists a stable topology.*

Proof. For $m < d$, the following topology is stable: Miners connect in a clique thus all have score $\frac{m-1}{m}$. All non-miners connect to all miners, either via incoming or outgoing edges. With global ordering, we have that in the miner clique, a miner with k outgoing edges to other miners has $d-1-k$ outgoing edges to the first $d-1-k$ non-miners by the global order. Every non-miner connects to all k miners it is not connected to and similarly has its remaining $d-1-k$ outgoing edges to the first $d-1-k$ (or if no global ordering, to any) of the remaining non-miners. Thus all non-miners have the same score, and no non-miner will drop an edge to a miner. We then set everyone's d-th edge, note miners never have a d-th edge since they are already connected to everyone. For non-miners, any new edge will be to a node not in the first $d-1-k$ non-miners defined above (or any non-miner for LIFO), and therefore will never win the tie-breaker.

For $d \leq m < 2d$, we present a stable topology with all miners connecting in a clique using $(m-1)/2$ out-going edges and all remaining outgoing edges from the miners going to some (top-ranked) n' non-miners who have a score of 1 and all outgoing edges into the clique. Thus the miners have a miner-stable topology and the n' non-miners will keep their edges to the miners as long as the miners keep their edges to these non-miners and vice versa. All non-miners that are not part of the n' non-miners connect to the top $d-1$

[5] Note that this is true of Bitcoin (where $d = 8$ and $n > 10000$) and Ethereum ($d \approx 16$ and $n > 2000$).

miners, thus the top $d-1$ miners have no d-th edge and the rest have an edge to some non-miner that will at best have a score of 1 but be higher-ranked than the n' non-miners. The m miners use $(m-1)/2$ outgoing edges each to connect to the clique, leaving $d-1-(m-1)/2$ to connect to the n' non-miners. The n' non-miners have a cumulative $n'(d-1)$ edges, that need to be enough to connect them to all m miners, thus $n'm = n'(d-1) + d - 1 - (m-1)/2$, yielding $n' = \frac{(2d-1)-m}{2}$. Note that when $m = 2d-1$ the miners use all their outgoing edges to connect in the clique so $n' = 0$. □

Next, we show how all remaining cases cannot reach stability. We provide a proof sketch, referring the reader to the full version [17] for the full proofs.

Lemma 6 (LIFO and Global Ordering Tie-Breaking Instability for Large m). *For all $m \geq 2d$, with Tie Breaker Rules LIFO or global ordering, there exists no stable topology for $n \gg d$.*

Proof. We essentially show that in a stable topology, 1.) there is a subset of miners who have the minimum score and *all nodes* must have all of their outgoing edges to this subset for the edge to be stable, and 2.) this subset does not have enough miners to ensure that all their outgoing edges are to one another, so some

Lemma 7 (FIFO and Random Tie-Breaking Instability for Large n). *For all m, d, and $n > 2d^2$ with Tie Breaker Rules FIFO or random, there exists no stable topology.*

Proof. Where miner-stable topologies exist, we show that even if all non-miners connect to all miners (i.e., have the same minimal score), any edge between themselves cannot be stable as the Tie Breaker Rule allows the new edge to replace an existing edge with the same score. □

All examples given in the affirmative proofs of Table 1 rely on specific topologies involving a highly connected sub-network of miners, and generally low-diameter networks. Next we show that such properties are in fact necessary for any stable network.

Properties of stable topologies We can extend the above analysis to derive the following central properties that *any* miner-stable or network-stable topologies must have. Principally, these state that stable networks must necessarily have a highly-connected miner component.

Proposition 3. *If $m < 2d$, and ties are broken by FIFO or randomly, the only miner-stable topology is a clique.*

Proof. The proof follows similarly to that of Proposition 2. Assume there is some non-clique miner-stable topology. Thus there is some miner m_i who is not connected to all miners. There is some probability that in the next round, all miners connect to this miner with at least one miner, m_j having a new out-going edge to m_i. Miner m_i now has score $\frac{m-1}{m}$, thus m_j will not drop the edge to m_i since it is either smaller than some existing edge or ties with all edges and by FIFO (or random) Tie Breaker Rule replaces some edge (has a probability of replacing some edge). □

Proposition 4. *For all m, d and a global order (or LIFO) Tie Breaker Rule, in any miner-stable topology, all miners must connect to the first (or some) $d-1$ miners.*

Proof. (Case 1: global ordering) Assume there is some miner-stable topology where there is a miner m_i who does not connect to some miner m_j where m_j is in the top $d-1$ miners. In any round, there is a chance that all miners connect to miner m_j with miner m_i forming an outgoing edge to m_j. Thus in this round m_j has the minimum score and won't be dropped by m_i since it either has a lower score than some outgoing edge of m_i or is ranked higher.

(Case 2: LIFO) Assume that there are less than $d-1$ miners who all miners connect to in a miner stable topology. There is some miner m_i with score $> \frac{m-1}{m}$ which some miner, m_j, is not connected to. There is a chance that in the next round all miners who are not yet connected to m_i, connected to it, including m_j with an outgoing edge to m_i. In this round, m_i has the minimum score of $\frac{m-1}{m}$, and since less than $d-1$ miners have such a score, m_i will now be in the top $d-1$ edges of m_j. □

Proposition 5. *Any miner-stable topology (regardless of m, d) must have miner diameter at most 2 and will lead to a network diameter of at most 3.*

Proof. This follows directly from requirements of stability. Assume there is some miner-stable topology where no miner connects to all miners. There is always some possibility that in a round all miners connect to some miner m_i with some miner m_j having a new outgoing edge to m_i. As has been argued, this new edge contradicts that the initial topology was miner-stable. Thus any stable topology has at least one miner who connects to all other miners. Since the lowest-scoring node is a miner, eventually all non-miners will connect to at least one miner and any future state has all non-miners connecting to at least one miner. Note there is no guarantee of stability of these edges, but we get that the distance between any two non-miners is at most 3. □

Recall from Lemma 1 that the Nash Equilibrium we present for the idealized game also has a miner diameter of 2 and network diameter of 3.

Inbound connection cap. We briefly explore the impact of inbound connection caps on the stability properties we explored above. Without loss of generality, we assume that the inbound cap $d_{in} > d$, otherwise nodes would not have the inbound capacity to form new connections.

Lemma 8. *For $m < 2d$ and any Tie Breaker Rule, there exists a miner-stable topology. For tie-breaking rules LIFO and global ordering, miner-stable topologies exist for $m \leq max(2d-1, d_{in} + d/2)$.*

Proof. For $m < 2d$, there are enough edges between miners for cliques to form, thus they are still stable under Proposition 1. For $m \geq 2d$, consider a topology where the first $d-1$ miners form a clique, call them set M, then the rest of the miners have their outgoing edges to all miners in M. Thus all miners in M have the minimum score, and won't be dropped by any miner; all miners not in M have their outbound capacity filled by these stable edges so they will not connect to each other. We now calculate the maximum miners not in M. The maximum in-bound capacity for all miners in M is for each miner to have equal incoming/outgoing edges $= \frac{d-2}{2}$. Thus they each have $d_{in} - \frac{d-2}{2}$ capacity left for miners not in M, this happens for $m \leq d - 1 + d_{in} - \frac{d-2}{2} = d_{in} + d/2$. □

It remains an open problem if there exists a miner-stable topology for other regimes.

4.3 Reachable Stability

In the previous section, we characterized the exact conditions for when miner-stable and network-stable topologies exist for the unbounded Single-Exploratory-Greedy protocol under different Tie Breaker Rules. We next consider *if* these stable topologies are necessarily reachable. We show that *for $m \leq d$, all Tie Breaker Rules will lead to miner-stable topologies*, in particular a miner-clique. Additionally, for $m < d$, we show that *if there exists a stable network, then the network will stabilize.* For larger values of m, our proofs of stability in the previous subsection lead us to believe that the stable networks that do exist are hard to reach, we leave this for future work.

We consider the state transition model of the Single-Exploratory-Greedy game. We define a state of the protocol as the graph G of the network at some round r, and the set of reachable states as the states possible after the execution of one round of the protocol.

Proposition 6 (Eventual Miner-Stable Clique). *If the number of miners $m \leq d$, the network will reach a miner-stable topology.*

Proof. To prove this we show that some miner-stable topology is reachable from any state (network topology and round). To do this we present the following *clique process* which defines an exact path from *any* state to some miner clique which we've shown in the previous section is the only miner-stable topology for $m \leq d$. For convenience, we define some arbitrary order over all nodes such that the miners are the first m nodes.

Def. *clique-process*:

For each round r, take $\min_i$ s.t. $\text{score}(m_i) > \frac{m-1}{m}$:

1. $\forall j \leq m$, if $e_{i,j} = 0$, then m_j chooses m_i as their new random outgoing edge.
2. $\forall j \leq i$ $\text{score}(m_j) = \frac{m-1}{m}$
3. Note: $\forall j \leq i$ no edges *to* m_j are dropped

For the remainder of the proof, we re-define an edge as stable in some state, if the value of the edge does not change for all future states on the path defined by the clique-process. Equally, a miner is miner-stable if all of its edges to other miners are stable in the set of future states defined by the clique-process path. Note that once all miners are miner-stable on the clique-process path, then the network is miner-stable for any path as this only happens when a clique is formed by the miners.

We now prove the following statement: Starting from any state, the path defined by the clique-process will result in the first j miners being miner-stable, for $j \leq m$.

Base case: (Miner m_1 will miner-stabilize) Assume that m_1 never stabilizes. This means $\exists j \leq m$ s.t. $e_{1,j}$ does not stabilize. By definition of the clique-process, $i \geq 1$ for all rounds, thus at step (2.) of each round score(m_1)=$\frac{m-1}{m}$, thus $e_{1,j} = 1$. For this edge to not be stable, $\exists r' \geq r$ where $e_{1,j} = 0$ at step (3.) which only happens if this edge was an outgoing edge from m_1, and m_1 drops the edge at round r'. Then, at round $r' + 1$ we have $i = 1$, so miner m_j will form a *stable* outgoing edge to m_1 (note step 3 of the clique-process).

Inductive step: Assume that $\forall j' < j \leq m$ miner $m_{j'}$ is miner stable at some round r. It must be that $\forall j' < j$ score($m_{j'}$) $= \frac{m-1}{m}$. Note that this means for each round $r' \geq r$, score(m_j) $= \frac{m-1}{m}$ at step (2.), even if m_j is not stable, as it must be that $i \geq j$. The proof that m_j stabilizes follows similarly to the base case.

Assume m_j never stabilizes, this means $\exists h \leq m$ s.t. $e_{j,h}$ does not stabilize. By definition of the clique-process, $i \geq j$ for all rounds, thus at step (2.) of each round score(m_j)=$\frac{m-1}{m}$, thus $e_{j,h} = 1$. For this edge to not be stable, $\exists r' \geq r$ where $e_{j,h} = 0$ at step (3.) which only happens if this edge was an outgoing edge from m_j, and m_j drops the edge at round r'. If this happens, at round $r' + 1$ we have that $i = j$, so miner m_h will form a *stable* outgoing edge to m_j (note step 3 of the clique-process). □

We now extend our results to include network stability guarantees. Recall that only Tie Breaker Rule rules LIFO and Global-ordering have stable networks and only for $m < 2d$. Note that the following proposition is for m strictly less than d as the stable networks we used in our proof of Lemma 5 for $d \leq m < 2d$ were quite contrived and may not be reachable.

Proposition 7 (Eventual Network Stability). *If the number of miners $m < d$, given Tie Breaker Rule of Global-ordering or LIFO, the network will arrive at a stable topology.*

Proof. We give a proof sketch. By Proposition 6, the miners will eventually form a miner-stable topology which is a miner clique. Given this, we first argue that all edges involving miners will stabilize, and lastly that all edges between non-miners will stabilize. To do so, we define similar stabilizing processes as in the proof above. For details, see [17].

Corollary 1. *For $m < d$, and any Tie Breaker Rule, if there exists a miner-stable or network-stable topology, then the network will arrive at some such topology.*

5 Simulations

In the previous section, we prove the conditions for stability in a network implementing Single-Exploratory-Greedy and some conditions for when a network running Single-Exploratory-Greedy will stabilize. In this section, we simulate Single-Exploratory-Greedy to explore how the networks evolve, possible properties that stabilize for different network sizes and in-bound connection capacities, and the impact of having more exploratory edges in the protocol.

As in the theory, we assume each node can uniformly pick a random new peer from the set of all nodes in the network, and knows the score of their peers at step (3) of each round. In simulations, at step (1) of the protocol, each node adds a new edge to a peer they are not yet connected to. The order of the peer selection is randomized in each round, and with capped d_{in}, if the new peer they try to connect to has full in-edge capacity, then the node tries another random node until they succeed, or run out of peers (e.g., if all nodes that are not currently peers have full in-connections)[6]. In the previous section, we prove that of the Tie Breaker Rules we consider, only LIFO and a global ordering have the property that for all m, n there exists a miner-stable topology. Since a global ordering over all nodes is not intuitive in a decentralized network, we

[6] We emphasize that all other steps of each round are done in parallel for all nodes. We explored simulations where nodes completed each full round one at a time and found no differences in the simulation results.

use LIFO as the Tie Breaker Rule in our simulations. We run each of the simulations below 20 times and plot the averages of the runs.

Main results of the simulations:

1. Our simulations suggest that the effect of the d_{in} cap is smooth (as the cap is raised, the behavior shifts smoothly to the no cap model).
2. Though theoretical results show the existence of stable topologies, in particular for $m \leq d$ that cliques between miners are stable and inevitable, even for low n (e.g., 400 and 900), arriving at such a topology is hard/improbable and our simulations do not reach a miner clique in the first 256 rounds.
3. An extension of Single-Exploratory-Greedy where nodes drop the k-worst outgoing edges each round suggests the additional randomness can lead to the network converging quicker to the same values, but that this advantage is maximized at low values of k (e.g., 3).

Our analysis indicates that, consistent with the theoretical results, a hierarchical structure emerges within these networks with respect to node scores, degree, and centrality, which are highly correlated. In general, these networks converge to overall lower properties, but often with some nodes (primarily the miners) being at a clear advantage and skewing the average.

Starting topology. For most of our simulations, we begin with a random graph where each node chooses d initial random peers it is not yet connected to (see Figure 2(a).). As we discuss below, the Single-Exploratory-Greedy protocol tends to converge to a structure like Figure 2(b) where few nodes in the middle connect to most others on the periphery. We also run simulations starting with a small-world graph (using the Watts-Strogratz [35] model with probability 0.5) and with a scale-free graph (using the Barabasi-Albert [1] model with an initial connected component of size 20). For a full analysis of these, see [17]. We observe that the initial graph does have a small impact on the properties the *capped* simulations are converging to (e.g., scale-free network converges to higher average distance to miners).

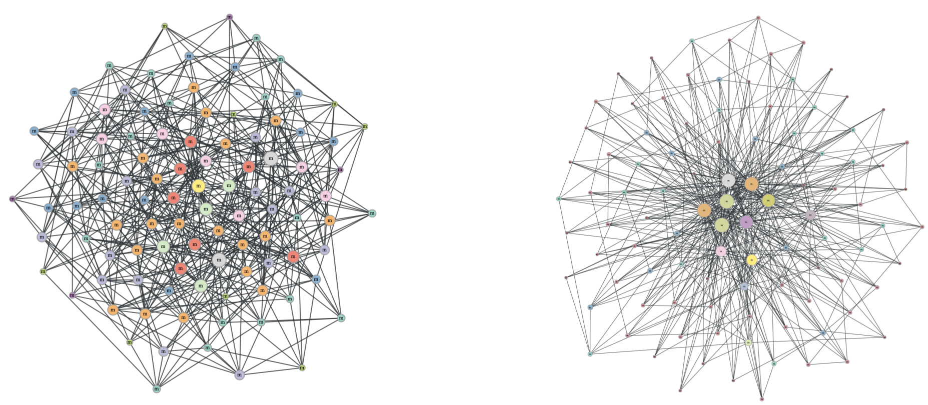

(a) (b)

Fig. 2: (a) the random graph created from all nodes choosing d random nodes to connect to. (b) the graph we seem to converge to from running Single-Exploratory-Greedy with all nodes being homogeneous miners, size scaled by degree. This structure appears after ∼20 rounds.

5.1 Effect of In-Degree Caps

To complement our limited theoretical analysis with bounded d_{in}, we explore the impact the inbound cap has on networks of all miners, and networks with a subset of 10 nodes being miners. Figure 3 shows the average distance over time for different inbound caps. We see that if the network is all miners, we need a substantial inbound cap for the average distance between all nodes to be better than the random graph within the span of our simulations. We see, however, in Figure 4 that this lowered average distance is skewed by a few nodes being better off than the majority of the network. When only some of the nodes are miners, the network has the potential to converge to one with a better average distance to the miners; higher the inbound cap, lower the average distance to miners. Figure 5 supports this, showing the network diameter approaching 2. Additionally, the diameter between just the miners is approaching 1 meaning that even with bounded capacity, the miners are forming a clique. *Our main observation here is that the value of the bound has a smooth effect on network properties.*

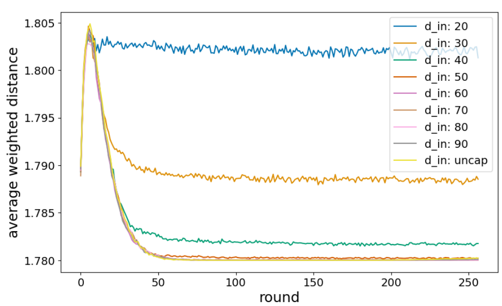

(a) All Miners

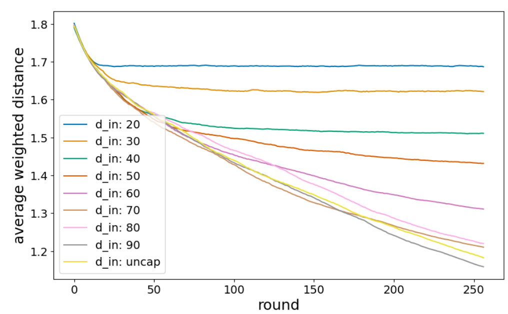

(b) Subset Miners

Fig. 3: Average distance to miners per round for an $n = 100$ node network with all miners and $m = 10$ miners, for different in-degree caps.

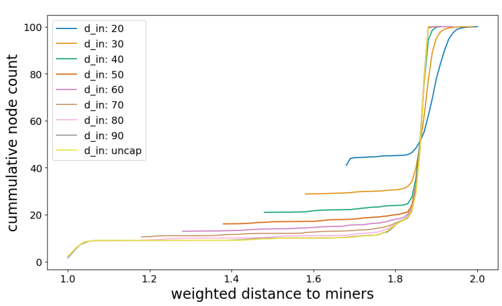

(a) All Miners

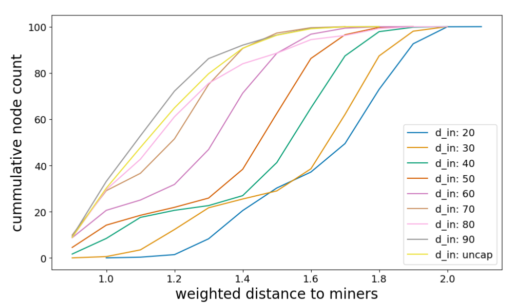

(b) Subset Miners

Fig. 4: Distribution of distance to miners(node scores) at round 256 for an $n = 100$ node network with all miners and $m = 10$ miners, for different in-degree caps.

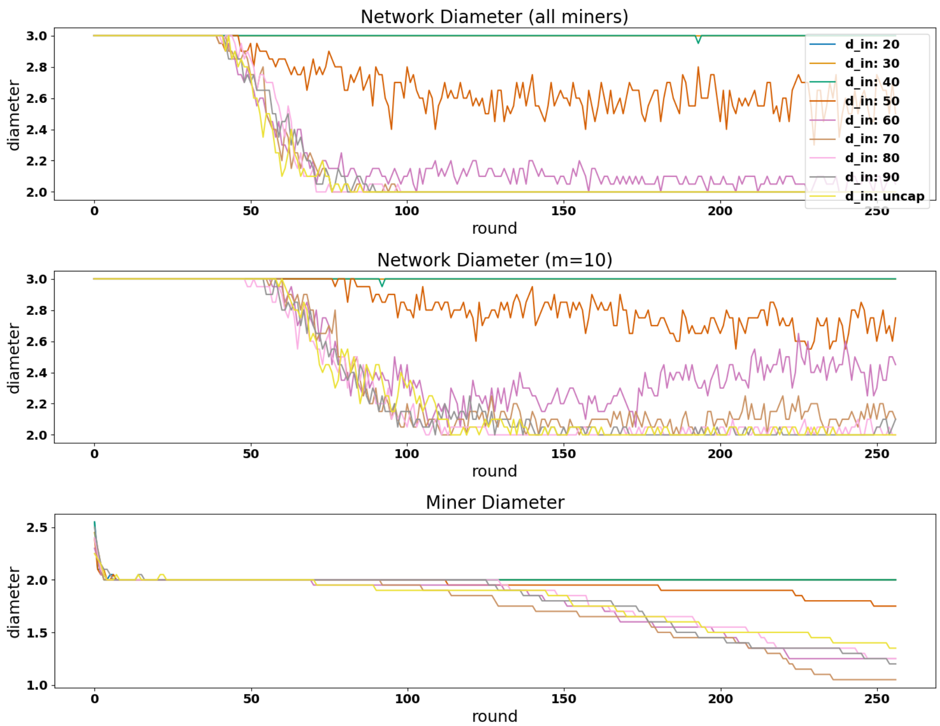

Fig. 5: Diameter of a 100 node network with all miners and a subset of 10 miners, and the diameter between the 10 miners.

5.2 Network Size

We also explore the impact of network size on our simulations when a subset of the network is miners. Again, we consider $m = 10$ and $d = 10$, with $d_{in} = 20$ or uncapped. Our theoretical results show that for $m \leq d$, the miners connecting in a clique is a miner-stable topology, regardless of what the rest of the network is doing (for uncapped d_{in}, this is the only miner-stable topology). We run simulations for $n = 100, 400$ and 900 and look at the network/miner diameter and eccentricities in [17]. *Our main take away is that, though miners are becoming better connected over time, with more non-miners, it is harder for the miners to form a clique.* This is likely due a decreasing probability of another miner being randomly selected by a miner as a neighbor as the number of nodes increases. The non-miners also maintain connections to the miners (seen in the eccentricity of the miners in [17]), it is thus possible the non-miners will occupy all the incoming edges of the miners before they form a clique. This seems to be the case for even $n = 100$ as the capped miner diameter is stabilizing at 2.

5.3 k-worst edges

So far, in both the theory and our simulations, we've considered dropping only the single worst outgoing edge. We also explore what happens when nodes drop the $k < d$ worst outgoing edges. Note that $k = d$ is essentially a random graph. Conceptually, more dropped edges leads to more randomness in the network, as more nodes have more incoming connections (see Figure 7). When all nodes are miners, there is not a big difference in network behavior (Figure 6 a. and b. show some difference in behavior but the scale is quite small). For most cases (low enough k) in the subset miner simulations, miners are finding each other over time, but the additional randomness in the network appears to lead to worse average distances (Figure 6 d.). The main advantage that higher values of k sometimes give is a faster convergence to the same values, we show this in [17]. *Generally, dropping more outgoing edges does not change the behavior of the network much, and where it does, the optimal k seems to be up to 3.*

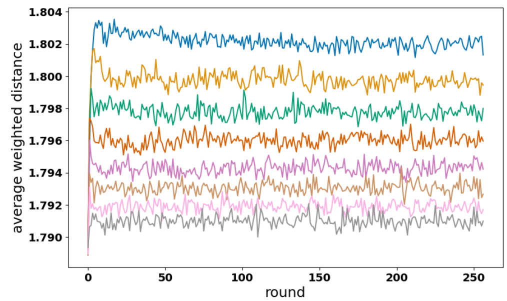

(a) all miners capped

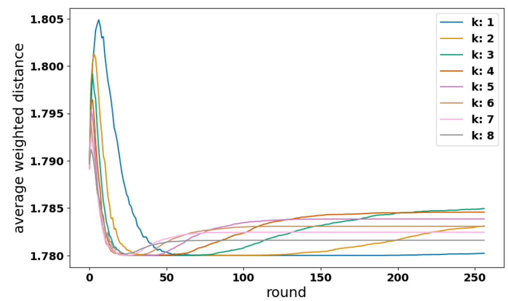

(b) all miners uncapped

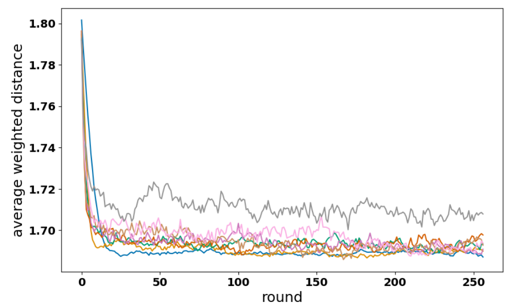

(c) 10 miners capped

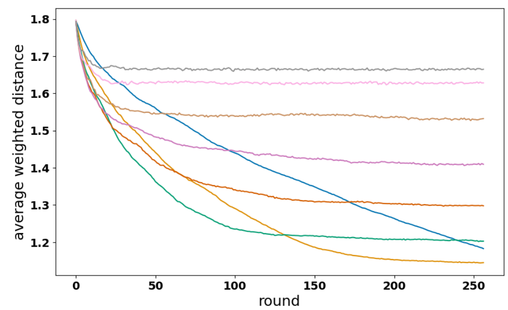

(d) 10 miners uncapped

Fig. 6: Average distance to miners for different k-worst values, for networks of $n = 100$, incoming cap of 20 or unlimited, and all miners or subset miners.

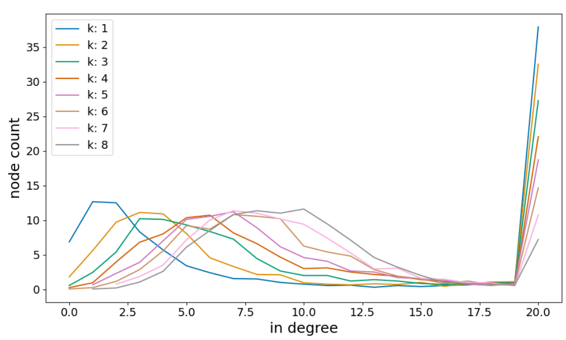

(a) capped

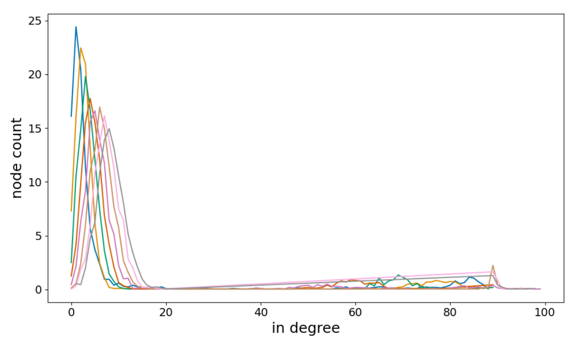

(b) uncapped

Fig. 7: In-degree distribution for different k-worst values for sub-network of miners. The case when all nodes are miners looks the same. For the capped case, the larger the k, the closer the network behaves to the initial random graph.

6 Discussion

In this work, we study the impact of greedy peering choices in a P2P network. In our models, we consider some nodes (either all or a subset of nodes) as "miners" as the source of messages that all nodes are trying to optimize their distance to. We mostly consider the case of homogeneous miners[7], so it is often the case that a nodes' peers will tie in their scores. In fact, in our theoretical results, the choice of tie-breaking rule is a determinant in whether the network will stabilize. Our choice of edges having a weight of 1 symbolizes the

[7] We explore miners having different weights (representing different mining powers) in [17] and find that the more disparate the weights, the more the network behaves like the subset miner network.

delay in propagation each hop adds, normalized over some period of time. In reality, this weight would be an average of delays, and therefore ties between peers would be quite rare. We consider our model of lowering hop distance (as opposed to link-latency distance) for a network where link latency has less of an impact than the processing latency added by each node. It seems that the global-ordering rule fits best as a stand-in for this, assuming the processing latency of any two nodes don't differ too much. Node heterogeneity of this kind and in default parameters is left for future work.

Both our theoretical results and simulations look at the conditions for stability and how the network evolves into a possibly stable state. In both cases, we see a two-tier system for the miners in a stable topology and an additional third tier of all other nodes. An interesting direction for future work is what happens when new nodes (with or without the preferential property) join these stable networks: Does the network (re)stabilize and how long does it take?

In both our simulations and theoretical analysis, stability properties are observed in the regime of a small m relative to n, we note that cryptocurrency networks historically fall under this regime with a minority of network participants being miners [18]. In our simulations, we explore how networks stabilize, particularly for network conditions that are difficult to theorize about. While we explore several parameter choices, one caveat of our simulations is that we stay within relatively small network sizes (compared to the major existing cryptocurrencies, e.g., Bitcoin and Ethereum) so that the computation is tractable and we are still able to observe network dynamics[8]. In the same spirit, we consider a single d value while changing the network size and d_{in} caps to explore the proportional effect of the out-degree. Our theoretical results, however, deal with more general parameters.

References

1. Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.
2. Vijeth Aradhya, Seth Gilbert, and Aquinas Hobor. Overchain: Building a robust overlay with a blockchain. *arXiv preprint arXiv:2201.12809*, 2022.
3. Kushal Babel and Lucas Baker. Strategic peer selection using transaction value and latency. In *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security*, pages 9–14, 2022.
4. Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.
5. Bitcoin core 24.0.1. `https://github.com/bitcoin/bitcoin`, 2023.
6. Today's cryptocurrency prices by market cap. `https://coinmarketcap.com/`, 2023.
7. Jacomo Corbo and David C. Parkes. The price of selfish behavior in bilateral network formation. In *Proc. of PODC'05*, pages 99–107, 2005.
8. Sergi Delgado-Segura, Surya Bakshi, Cristina Pérez-Solà, James Litton, Andrew Pachulski, Andrew Miller, and Bobby Bhattacharjee. Txprobe: Discovering Bitcoin's network topology using orphan transactions. In *Financial Cryptography and Data Security: 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18–22, 2019, Revised Selected Papers 23*, pages 550–566. Springer, 2019.
9. Erik D. Demaine, Mohammadtaghi Hajiaghayi, Hamid Mahini, and Morteza Zadimoghaddam. The price of anarchy in network creation games. *ACM Trans. Algorithms*, 8(2):1–13, 2012.
10. devp2p. `https://github.com/ethereum/devp2p`, 2023.
11. Maya Dotan, Yvonne-Anne Pignolet, Stefan Schmid, Saar Tochner, and Aviv Zohar. Survey on blockchain networking: Context, state-of-the-art, challenges. *ACM Computing Surveys (CSUR)*, 54(5):1–34, 2021.
12. Matthieu Gouel, Kevin Vermeulen, Olivier Fourmaux, Timur Friedman, and Robert Beverly. IP geolocation database stability and implications for network research. In *Network Traffic Measurement and Analysis Conference*, 2021.
13. Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. Eclipse attacks on Bitcoin's peer-to-peer network. In *24th {USENIX} Security Symposium ({USENIX} Security 15)*, pages 129–144, 2015.
14. Sebastian Henningsen, Daniel Teunis, Martin Florian, and Björn Scheuermann. Eclipsing Ethereum peers with false friends. *arXiv preprint arXiv:1908.10141*, 2019.
15. Matthew Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74, 1996.

[8] We note, however, that our simulations are within the realm of many smaller cryptocurrency networks such as the Ethereum Classic Network with around 500 observed nodes [28].

16. Lucianna Kiffer, Joachim Neu, Srivatsan Sridhar, Aviv Zohar, and David Tse. Security of blockchains at capacity. *arXiv preprint arXiv:2303.09113*, 2023.
17. Lucianna Kiffer and Rajmohan Rajaraman. Stability of p2p networks under greedy peering (full version). http://arxiv.org/abs/2402.14666, 2024.
18. Lucianna Kiffer, Asad Salman, Dave Levin, Alan Mislove, and Cristina Nita-Rotaru. Under the hood of the Ethereum gossip protocol. In *International Conference on Financial Cryptography and Data Security*, pages 437–456. Springer, 2021.
19. Nikolaos Laoutaris, Laura Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng. Bounded budget connection (BBC) games or how to make friends and influence people, on a budget. *Journal of Computer and System Sciences*, 80(7):1266–1284, 2014.
20. Kai Li, Yuzhe Tang, Jiaqi Chen, Yibo Wang, and Xianghong Liu. Toposhot: uncovering Ethereum's network topology leveraging replacement transactions. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 302–319, 2021.
21. Lighthouse book: Advanced networking. https://lighthouse-book.sigmaprime.io/advanced_networking.html.
22. Cristian Lumezanu, Randy Baden, Neil Spring, and Bobby Bhattacharjee. Triangle inequality variations in the internet. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*, pages 177–183, 2009.
23. Yifan Mao, Soubhik Deb, Shaileshh Bojja Venkatakrishnan, Sreeram Kannan, and Kannan Srinivasan. Perigee: Efficient peer-to-peer network design for blockchains. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 428–437, 2020.
24. Yuval Marcus, Ethan Heilman, and Sharon Goldberg. Low-resource eclipse attacks on Ethereum's peer-to-peer network. *Cryptology ePrint Archive*, 2018.
25. Eli A. Meirom, Shie Mannor, and Ariel Orda. Network formation games with heterogeneous players and the internet structure. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, page 735–752, New York, NY, USA, 2014. Association for Computing Machinery.
26. Andrew Miller, James Litton, Andrew Pachulski, Neal Gupta, Dave Levin, Neil Spring, and Bobby Bhattacharjee. Discovering Bitcoin's public topology and influential nodes. *et al*, 2015.
27. Joachim Neu, Srivatsan Sridhar, Lei Yang, David Tse, and Mohammad Alizadeh. Longest chain consensus under bandwidth constraint. In *AFT*. ACM, 2022.
28. Etc node explorer. https://etcnodes.org/. Accessed on 08-08-2023.
29. M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
30. Sehyun Park, Seongwon Im, Youhwan Seol, and Jeongyeup Paek. Nodes in the Bitcoin network: Comparative measurement study and survey. *IEEE ACCESS*, 7:57009–57022, 2019.
31. Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.
32. Elias Rohrer and Florian Tschorsch. Kadcast: A structured approach to broadcast in blockchain networks. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 199–213, 2019.
33. Weizhao Tang, Lucianna Kiffer, Giulia Fanti, and Ari Juels. Strategic latency reduction in blockchain peer-to-peer networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(2):1–33, 2023.
34. Bhavesh Toshniwal and Kotaro Kataoka. Comparative performance analysis of underlying network topologies for blockchain. In *2021 International Conference on Information Networking (ICOIN)*, pages 367–372. IEEE, 2021.
35. Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.
36. Jan Zich, Yoshiharu Kohayakawa, Vojtech Rödl, and Vaidy Sunderam. Jumpnet: Improving connectivity and robustness in unstructured P2P networks by randomness. *Internet Mathematics*, 5(3):227–250, 2008.