

Covert Networks: How Hard is It to Hide?

Palash Dey
 Indian Institute of Technology
 Kharagpur
 palash.dey@cse.iitkgp.ac.in

Sourav Medya
 University of California
 Santa Barbara
 medya@cs.ucsb.edu

ABSTRACT

Covert networks are social networks that often consist of harmful users. Social Network Analysis (SNA) has played an important role in reducing criminal activities (e.g., counter terrorism) via detecting the influential users in such networks. There are various popular measures to quantify how influential or central any vertex is in a network. As expected, strategic and influential miscreants in covert networks would try to hide herself and her partners (called *leaders*) from being detected via these measures by introducing new edges.

Waniek et al. [72] show that the corresponding computational problem, called Hiding Leader, is NP-complete for the degree and closeness centrality measures. We study the popular core centrality measure and show that the problem is NP-complete even when the core centrality of every leader is only 3. On the contrary, we prove that the problem becomes polynomial time solvable for the degree centrality measure if the degree of every leader is bounded above by any constant. We then focus on the optimization version of the problem and show that the Hiding Leader problem admits a 2 factor approximation algorithm for the degree centrality measure. We complement it by proving that one cannot hope to have any $(2 - \epsilon)$ factor approximation algorithm for any constant $\epsilon > 0$ unless there is a $\epsilon/2$ factor polynomial time algorithm for the Densest k -Subgraph problem which would be considered a significant breakthrough. We empirically establish that our 2 factor approximation algorithm frequently finds out a near optimal solution. On the contrary, for the core centrality measure, we show that the Hiding Leader problem does not admit any $(1 - \alpha) \ln n$ factor approximation algorithm for any constant $\alpha \in (0, 1)$ unless $P = NP$ even when the core centrality of every leader is only 3. Hence, our work shows that, although classical complexity theoretic framework fails to shed any light on relative difficulty of Hiding Leader for different centrality measures, the problem is significantly “harder” for the core centrality measure than the degree centrality one.

KEYWORDS

Social Network Analysis; Centrality; Hiding in Networks; Algorithms; Social Influence; Covert Network; Core; Approximation Algorithm; Network Design; Densest Subgraph

ACM Reference Format:

Palash Dey and Sourav Medya. 2019. Covert Networks: How Hard is It to Hide?. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 10 pages.

Dey is funded by DST INSPIRE grant no. 04/2016/001479 and IIT Kharagpur grant no. IIT/SRIC/CS/VTS/2018-19/247.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Social network analysis (SNA) has played a pivotal role in many applications in multi-agent systems and artificial intelligence [16, 55, 63, 71]. One of the most successful applications of SNA is in counter-terrorism via analyzing *covert networks* [17, 45, 60, 75]. Covert network loosely refers to network of criminals, terrorists, illegal activities, etc. Security personnel regularly use various SNA tools to understand criminal behavior, catch their leaders, and effectively dismantle such networks [27, 30, 40].

Centrality measure is one of the most useful tools that SNA provides to analyze covert networks. It assigns scores to the vertices based on their relative *influence or importance* in the network [8]; depending on the centrality measure, higher scores may correspond to important vertices and important vertices are expected to be more central. One of the simplest and oldest such centrality measures is the *degree centrality* which ranks vertices according to their degree [67]. Other important examples include *closeness centrality* and *betweenness centrality* that are measures based on shortest paths [9]. Another centrality measure is the *core centrality* [65] which ranks the vertices based on their *core number*. Intuitively speaking, if a vertex has a high core number, then it is part of some dense cohesive community within the network. Formally, a k -core is an induced subgraph of the network where the minimum degree of the vertices is at least k . The core number of a vertex is the highest integer k such that the vertex is part of some k -core. Therefore, the core centrality can be more revealing about the position of a node than its degree centrality—while degree centrality only concerns about the degree of a vertex, the core centrality elegantly takes into consideration the degrees of the neighbors as well as the vertex. These two measures are also related in the sense that the core centrality of any vertex is at most its degree centrality. Due to its sophisticated nature, the core centrality has been extensively used in the study of covert networks [51, 53, 66] as well as in other important tasks such as viral marketing and social engagement [12, 39] in social networks.

In this paper, our goal is to study the centrality measure based secrecy in covert networks. Indeed, understanding covert networks remains a challenging task mainly due to incompleteness and dynamic evolution of the data as well as the strategic nature of the users [1, 5, 41, 62, 64, 68]. Since the criminals often possess technical expertise [4, 15, 19, 36, 69], we are interested in the evolution of terrorist networks under a framework of strategic users [72]: *How is the network designed to hide the central or influential users aka the leaders?*

Waniek et al. [72] first propose the Hiding Leader problem which incorporates the viewpoint of the leaders of a criminal organization. It also explicitly models knowledge of the criminals about SNA tools that are used to detect them and thus help in dismantling their organization. Intuitively, the input in the Hiding Leader

problem is a network with a subset of vertices marked as leaders. The goal is to add fewest edges to ensure that various SNA tools do not rank any leader high based on centrality measures thereby capturing the *efficiency vs secrecy dilemma* that the criminals are believed to possess [20, 53, 70]. Waniek et al. show promising results that the HIDING LEADER problem is computationally intractable even for the simplest degree centrality measure.

1.1 Contribution

In this paper, we study the HIDING LEADER problem for the core centrality measure and show that the degree centrality measure is much more computationally vulnerable than the core centrality measure although the HIDING LEADER problem is NP-complete for both of them. We reinforce our above claim further through extensive empirical evaluations. Our specific contribution in this paper are as follows.

- ▷ We show that the HIDING LEADER problem for degree centrality is polynomial time solvable if the degree of every leader is bounded by some constant [Theorem 3.1].
- ▷ We present a 2 factor approximation algorithm for the HIDING LEADER problem for degree centrality which optimizes the number of edges added [Theorem 3.2]. We complement this by proving that, if there exists a $(2 - \epsilon)$ factor approximation algorithm for the above problem for any constant $0 < \epsilon < 1$, then there exists a $\epsilon/2$ factor approximation algorithm for the DENSEST k -SUBGRAPH problem [Theorem 3.3] which would be considered a substantial breakthrough. To the best of our knowledge, the state of the art algorithm for the DENSEST k -SUBGRAPH problem achieves an approximation ratio of $\tilde{O}(n^{1/4})$ only [11].
- ▷ For the core centrality measure, we show that the HIDING LEADER problem is NP-complete even if the core centrality of every leader is exactly 3 [Theorem 4.1]. We prove that our result is almost tight in the sense that the HIDING LEADER problem is polynomial time solvable if the core centrality of every leader is at most 1 [Proposition 4.3]. Moreover, we also prove that there does not exist any $(1 - \alpha) \ln n$ factor approximation algorithm for any constant $\alpha \in (0, 1)$ which optimizes the number of edges that one needs to add even when the core centrality of every leader is 3 [Corollary 4.2].
- ▷ We show that a construction of a network by Waniek et al. [72], called “captain network” there, hides the leaders with respect to the core centrality measure also.
- ▷ We empirically evaluate our 2-approximation algorithm for the degree centrality measure in synthetic networks. We observe that our algorithm almost always produces near optimal results in practice. In the experimental results, we also show the extent in which a leader can hide in the captain network with respect to core centrality.

1.2 Related Work

Waniek et al. first proposed and studied the HIDING LEADER problem [72, 73]. They proved that the problem is NP-complete for

both the degree and closeness centrality measures. They also proposed a procedure to design a captain (covert) network from scratch which not only hides the leaders based on the degree, closeness, and betweenness centrality measures, but also keeps the influence of the leaders high in the network. In this paper, we provide two approximability results for degree centrality and core centrality respectively. We also show the problem is harder in the case of core centrality. Liu et al. [44] studied another related problem to make the degree of each node in the network beyond a given constant by adding minimal edges.

Other problems that align with privacy issues in social networks were studied before [2, 77]. In [77], the authors showed how an adversary exploits online social networks to find the private information about users. Altshuler et al. [2] discussed the threat of malware targeted at extracting information in a real-world social network.

Computing centrality and related problems. A significant amount of related work study the computational complexity of various centrality measures. Brandes [13] first proposed an efficient algorithm to compute the betweenness centrality of a vertex in a network. More recently, Riondato et al. [61] introduced an approach to compute the top- k vertices according to the betweenness centrality using VC-dimension theory. Yoshida [76] studied similar problems for both the betweenness and coverage centrality measures in a group setting. Mahmood et al. subsequently improved the performance of the above algorithms using a novel sampling scheme [46]. There is an active line of research to optimize the centrality of one node as well as of a set of nodes [18, 21, 34, 49]. Nikos et al. proposed a novel procedure to maximize the expected decrease in shortest path distances from a given node to the remaining nodes via edge addition [59]. Crescenzi et al. [18] proposed greedy algorithms to increase centrality of certain vertices and show effectiveness of their approach through extensive simulation. Kilberg [38] and others studied behavioral models to understand why certain network topologies are common in covert networks [10, 19, 23]. Enders and Su [29] and others develop models to explain various properties like efficiency vs secrecy dilemma etc. of covert networks [26, 28, 35, 43]. Other important direction includes quantifying the influence of vertices; most prominent among them include *Independent Cascade* model [33], *Linear Threshold* model [37], Bass model [7, 47], etc.

Other network design problems. We also provide a few details about previous work on other network modification (design) problems. A set of design problems were introduced in [56]. Lin et al. [42] addressed a shortest path optimization problem via improving edge weights on undirected graphs. The node version of this problem was also studied [25, 48, 50]. Meyerson et al. [52] proposed approximation algorithms for single-source and all-pair shortest paths minimization. Faster algorithms for some of these problems were also presented in [57, 58]. Demaine et al. [22] minimized the diameter of a network by adding shortcut edges. Dey et al. [24] studied the social network effect in the surprise in elections.

2 PRELIMINARIES

For a positive integer ℓ , we denote the set $\{1, 2, \dots, \ell\}$ by $[\ell]$. A network or graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a tuple consisting of a finite set \mathcal{V} (or $\mathcal{V}[\mathcal{G}]$) of n vertices and a set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ of edges (also denoted

by $\mathcal{E}[\mathcal{G}]$). A network is called *undirected* if we have $(x, y) \in \mathcal{E}$ whenever we have $(y, x) \in \mathcal{E}$ for any $x, y \in \mathcal{V}$ with $x \neq y$. A *self loop* is an edge of the form (x, x) for some $x \in \mathcal{V}$. In this paper, we focus on undirected networks without any self loop. The *degree* of a vertex x is the number of edges incident on it which is $|\{e \in \mathcal{E} : x \in e\}|$. A *subgraph* of a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a network $\mathcal{H} = (\mathcal{U}, \mathcal{F})$ such that $\mathcal{U} \subseteq \mathcal{V}$ and $\mathcal{F} \subseteq \mathcal{E} \cap (\mathcal{U} \times \mathcal{U})$. For a positive integer k , a subgraph \mathcal{H} of a network \mathcal{G} is called a *k-core* if the degree of every vertex in \mathcal{H} is at least k . The *core number* of a vertex x in a network is the largest integer k such that x belongs to a k -core.

2.1 Network Centrality

Let \mathcal{G} be any network. Bavelas [8] introduces the notion of centrality of vertices. Intuitively, centrality measures try to capture the importance of a vertex in a network. Shaw [67] proposes the *degree centrality* measure which has turned out to be one of the most useful measures. The degree centrality of a vertex x in \mathcal{G} is the degree $\deg_{\mathcal{G}}(x)$ of x in the network, that is $|\{y \in \mathcal{V}[\mathcal{G}] : \{x, y\} \in \mathcal{E}[\mathcal{G}]\}|$.

Seidman [65] introduces the idea of *core centrality* which is particularly useful for finding network cohesion. For an integer k , a *k-core* is a subgraph \mathcal{H} of \mathcal{G} such that the degree of every vertex in \mathcal{H} is at least k . The core number of a vertex x in \mathcal{G} is the largest k such that x belongs to a k -core, that is $\max\{k \in \mathbb{N} : \exists \mathcal{H} \subseteq \mathcal{G}, \deg_{\mathcal{H}}(x) \geq k \forall x \in \mathcal{V}[\mathcal{H}]\}$. The core centrality of a vertex x in \mathcal{G} is its core number in the network. Other popular network centrality measures includes closeness centrality [9], betweenness centrality [3, 31], etc.

2.2 Problem Definition

Intuitively, the input in the HIDING LEADER problem is a network with a subset of vertices marked as leaders (and the other vertices are followers), a budget b which is the maximum number of edges that we can add in the network, and a target d which is the minimum number of followers whose centrality must be at least as high as the centrality of any leader in the resulting network (after addition of the new edges). We now define our problem formally. In Definition 2.1, $c(\cdot, \cdot)$ denote either degree centrality or core centrality.

Definition 2.1 (HIDING LEADER (HL)). Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a subset $\mathcal{L} \subseteq \mathcal{V}$ of leader vertices, an integer b denoting the maximum number of edges that we are allowed to add in \mathcal{G} , an integer d denoting the number of follower vertices in $\mathcal{F} = \mathcal{V} \setminus \mathcal{L}$ whose final centrality should be at least as high as any leader, the goal is to compute if there exists a subset $\mathcal{W} \subseteq \mathcal{F} \times \mathcal{F}$ of edges between followers such that the conditions below hold.

- (i) $|\mathcal{W}| \leq b$
- (ii) $\exists \mathcal{F}' \subseteq \mathcal{F} |\mathcal{F}'| \geq d$ such that

$$c(\mathcal{G}', f) \geq c(\mathcal{G}', l), \forall f \in \mathcal{F}', l \in \mathcal{L}$$
 where $\mathcal{G}' = (\mathcal{V}, \mathcal{E} \cup \mathcal{W})$.

3 RESULTS FOR DEGREE CENTRALITY

We present our algorithmic and hardness results for the HIDING LEADER problem for the degree centrality measure in this section. We begin with presenting our polynomial time algorithm for the

HIDING LEADER problem for the degree centrality measure when the degree of every leader in the network is bounded above by any constant. On a high level, our algorithm makes greedy choices as long as it can and uses local search technique when “stuck.”

THEOREM 3.1. *There exists a polynomial time algorithm for the HIDING LEADER problem for degree centrality if the degree of every leader is bounded by any constant.*

PROOF. Let \mathcal{G} be the input graph and k the highest degree of any leader; that is $k = \max\{\deg_{\mathcal{G}}(l) : l \in \mathcal{L}\}$. We are given that k is a constant. If the number of followers is at most $2k$, then there are at most $\binom{2k}{2}$ (which is a constant) new edges that we can add and we try all possible subsets of it of cardinality at most b . The number of such subsets is at most $2^{\binom{2k}{2}}$ which is a constant and thus we can output correctly in polynomial time. So let us assume that the number of followers is at least $2k + 1$. Similarly we can also assume without loss of generality that the budget b is at least $4k^2$ since otherwise we will try to add all possible b new edges (there are only $O(n^{2b}) = n^{O(1)}$ possibilities since $k = O(1)$) and thus we can output correctly in polynomial time.

Suppose there are already d' number of followers in \mathcal{G} whose degrees are at least k . If $d' \geq d$, we output YES. Otherwise let us assume without loss of generality that $d' < k$. Let $\mathcal{X} \subseteq \mathcal{F}$ with $|\mathcal{X}| = d - d'$ be the set of top $d - d'$ highest degree followers in the network whose degrees are less than k . Intuitively, our algorithm greedily adds new edges between two vertices in \mathcal{X} whose degrees are less than k until it is stuck and removes some edge it had added before to make progress. Concretely, our algorithm works as follows. To distinguish existing (old) edges from newly added edges (by the algorithm) in \mathcal{G} , we color the existing (old) edges as red and whenever we add a new edge, we color it green. To begin with, all the edges in \mathcal{G} are colored red and there is no green edge. We apply the following step (\star) as long as we can. If the number of green edges in \mathcal{G} is less than b and there exist two vertices $x, y \in \mathcal{X}$ such that the degrees of both the vertices are less than k and there is no edge between them, then we add an edge $\{x, y\}$ in \mathcal{G} and color it green. Such a pair of vertices (if exists) can be found in $O(n^2)$. If such a pair of vertices does not exist in \mathcal{G} , then one of the following four cases must hold.

Case 1: The degree of every vertex in \mathcal{X} is at least k . In this case, we output YES.

Case 2: The number of green edges in \mathcal{G} is b . In this case, we output YES if the degree of every vertex in \mathcal{X} is at least k ; otherwise we output NO.

Case 3: There exists exactly one vertex $x \in \mathcal{X}$ with degree less than k . If the degree of x is $k - 1$, then we add an edge between x and any vertex $y \in \mathcal{L}$ such that there is no edge between x and y in \mathcal{G} and color it green. If the number of green edges in \mathcal{G} is at most b , then we output YES; otherwise we output NO. Otherwise we assume that the degree of x is less than $k - 1$. If the number of green edges in \mathcal{G} is at most $3k^2$, then we can add k new green edges on x and answer YES since we have already assumed that $b \geq 4k^2$. So let us assume without loss of generality that the number of green edges in \mathcal{G} is more than $3k^2$. Let $\{u, v\}$ be a green edge such that there is no edge between x and u and between x and v in \mathcal{G} . Such a green edge $\{u, v\}$ always exists in \mathcal{G} since the degree of x is less than $k - 1$

in \mathcal{G} and there are more than $3k^2$ green edges in \mathcal{G} . Moreover such an edge can be found in polynomial time by simply checking all the green edges. We now remove the green edge $\{u, v\}$ from \mathcal{G} , add two edges $\{x, u\}$ and $\{x, v\}$ in \mathcal{G} , color both of them green, and continue (return to the step (\star)).

Case 4: For every pair of vertices $x, y \in \mathcal{X}, x \neq y$ with degree less than k for both the vertices, there is an edge between them. Let $\mathcal{Z} \subseteq \mathcal{X}$ be the set of vertices in \mathcal{X} with degree less than k . Since there exists an edge between every pair of vertices in \mathcal{Z} in this case and the degree of every vertex in \mathcal{Z} is less than k , we have $|\mathcal{Z}| \leq k$. If the number of green edges in \mathcal{G} is at most $3k^2$, then we can add k new green edges on every vertex in \mathcal{Z} and answer YES since we have $|\mathcal{Z}| \leq k$ and $b \geq 4k^2$. So let us assume that the number of green edges in \mathcal{G} is more than $3k^2$. Let $a, b \in \mathcal{X}$ be any two vertices. Let $\mathcal{N}(a)$ and $\mathcal{N}(b)$ denote the set of neighbors of a and b in \mathcal{X} . Since the degrees of both a and b are less than k , we have $|\mathcal{N}(a)| < k$ and $|\mathcal{N}(b)| < k$. Since the degree of every vertex in \mathcal{X} is at most k and the number of green edges in \mathcal{G} is at least $3k^2$, there exists at least one green edge $\{u, v\}$ in \mathcal{G} which does not incident on any vertex in $\mathcal{N}(a) \cup \mathcal{N}(b) \cup \{a, b\}$ (there can be at most $2k^2$ green edges incident on any vertex in this set). Moreover such an edge can be found in polynomial time by simply checking all the green edges. We now remove the green edge $\{u, v\}$ from \mathcal{G} , add two edges $\{a, u\}$ and $\{b, v\}$ in \mathcal{G} , color both of them green, and continue (return to the step (\star)).

The algorithm always terminates in polynomial time since in every iteration, it adds a green edge and at most b ($< n^2$) green edges could be added – in cases 3 and 4, we have added two green edges and removed only one green edge; the algorithm terminates in cases 1 and 2. Also, whenever the algorithm outputs YES, adding green edges to the graph makes the degree of at least d followers in the network at least k . Hence, if the algorithm outputs YES, the instance is indeed a YES instance. So, let us assume that the algorithm outputs NO. Except (in case 3) when there exists exactly one vertex x in the network with degree less than k and the degree of x is $k - 1$, whenever we add one green edge in total (which is the same as adding two green edges and removing one green edge in cases 3 and 4), the sum of the degrees of all the vertices in \mathcal{X} increases by 2. Hence the number of green edges added in the graph is at most $\text{ALG} = \lceil \sum_{x \in \mathcal{X}} (k - \text{deg}(x)) / 2 \rceil$. Since the algorithm outputs NO, we have $\text{ALG} > b$. We observe that since any edge increases the degree of at most 2 vertices, when the algorithm outputs NO, the instance is indeed a NO instance. Hence the algorithm is correct. \square

We now present a simple 2 factor polynomial time approximation algorithm for the HIDING LEADER problem for degree centrality.

THEOREM 3.2. *There exists a polynomial time algorithm (HLDA) for approximating the budget b in HIDING LEADER within a factor of 2 for degree centrality.*

PROOF. Let $\mathcal{F}' \subseteq \mathcal{F}$ be the set of followers in \mathcal{F} whose degree centrality is at least the degree centrality of every vertex in \mathcal{L} . Let $|\mathcal{F}'| = d'$. If $d' \geq d$, then we output an empty set of edges. Let $x_i \in \mathcal{F}, i \in [d - d']$ be the $(d - d')$ followers with highest degree centrality among the vertices in $\mathcal{F} \setminus \mathcal{F}'$. We keep on adding edges with at least one end point in $\{x_i : i \in [d - d']\}$ until the degree centrality of every $x_i, i \in [d - d']$ is at least the degree of

every vertex in \mathcal{L} in the resulting graph. When we have d followers with degree at least the degree of every vertex in \mathcal{L} , we output the set of edges that we have added. The algorithm adds at most $\sum_{i=1}^{d-d'} (d - \text{deg}(x_i))$ many edges where $\text{deg}(x_i)$ is the degree of the vertex x_i in the input graph. Since any new edge can increase the sum of the degrees of the followers by at most 2, we have $\text{OPT} \geq \sum_{i=1}^{d-d'} (d - \text{deg}(x_i)) / 2 \geq \text{ALG} / 2$ by the choice of \mathcal{F}' . Hence our algorithm approximates b by a factor of 2. \square

We now complement our approximation algorithm in Theorem 3.2 by proving that if there exists a polynomial time approximation algorithm for the HIDING LEADER problem for degree centrality with approximation factor $(2 - \epsilon)$ for any constant $\epsilon > 0$, then there exists a constant factor polynomial time approximation algorithm for the DENSEST k -SUBGRAPH problem. In the DENSEST k -SUBGRAPH problem, the input is a graph \mathcal{G} and an integer k and we need to find a subgraph \mathcal{H} of \mathcal{G} on k vertices with highest density. The density of a graph on n vertices is the number of edges in it divided by $\binom{n}{2}$. To the best of our knowledge, we do not know whether there exists any polynomial time algorithm which can distinguish a graph containing a clique of size k from a graph where the density of every sub-graph of size k is at most $(\epsilon/2)$ (any $\epsilon/2$ factor approximation algorithm for the DENSEST k -SUBGRAPH problem would be able to distinguish). In fact, none of the known algorithms can distinguish even for some sub-constant values for ϵ (see [14] and references therein). We now show that if there exists a $(2 - \epsilon)$ factor approximation algorithm for the HIDING LEADER problem for any constant $0 < \epsilon < 1$, then there exists an $\epsilon/2$ factor approximation algorithm with the same running time (of the HIDING LEADER algorithm).

THEOREM 3.3. *Suppose there exists a $(2 - \epsilon)$ factor polynomial time approximation algorithm for the HIDING LEADER problem for degree centrality for some constant ϵ . Then there exists a polynomial time algorithm for distinguishing a graph containing a clique of size k from a graph where the density of every sub-graph of size k is at most $\epsilon/2$.*

PROOF. Let \mathcal{G} be any graph which satisfies either (exactly) one of the following properties.

- (i) **Completeness:** There exists a clique of size k in \mathcal{G} .
- (ii) **Soundness:** The density of any subgraph of \mathcal{G} of size k is at most $\epsilon/2$.

From \mathcal{G} we construct an instance of HIDING LEADER. Intuitively, we introduce a vertex a_v corresponding to every vertex $v \in \mathcal{V}[\mathcal{G}]$ in \mathcal{G} and the edge set among those vertices is the complement of the corresponding edge set in \mathcal{G} . To ensure that, in the resulting graph, the degree of every vertex a_v for $v \in \mathcal{V}[\mathcal{G}]$ is n , we add appropriate number of edges between a_v and some auxiliary vertices $d_{(v, \ell)}$ for every $v \in \mathcal{V}[\mathcal{G}], \ell \in [n]$; we ensure that the degree of any such auxiliary vertex is at most 1 which will guarantee that these auxiliary vertices are never part of any optimal solution. Finally we add a clique on a set $\{x_i : i \in [n + k]\}$ of leader vertices so that the degree centrality of every leader is $n + k - 1$. Formally, the instance

$(\mathcal{H}, \mathcal{L}, d)$ of HIDING LEADER is defined as follows.

$$\begin{aligned} \mathcal{V}[\mathcal{H}] &= \{a_v, d_{(v,\ell)} : v \in \mathcal{V}[\mathcal{G}], \ell \in [n]\} \cup \{x_i : i \in [n+k]\} \\ \mathcal{L} &= \{x_i : i \in [n+k]\} \\ \mathcal{E}[\mathcal{H}] &= \{\{a_u, a_v\} : (u, v) \notin \mathcal{E}[\mathcal{G}]\} \\ &\quad \cup \{\{x_i, x_j\} : 1 \leq i < j \leq n+k\} \\ &\quad \cup \{\{a_v, d_{(v,\ell)}\} : v \in \mathcal{V}[\mathcal{G}], \ell \in [\deg_{\mathcal{G}}(v) + 1]\} \\ d &= k \end{aligned}$$

We now use the $(2 - \varepsilon)$ factor approximation algorithm for the HIDING LEADER problem for degree centrality which outputs that there is a way to add b_{ALG} number of edges so that there exist at least d followers in \mathcal{H} whose degree is at least the degree of any leader. Let b_{OPT} denotes the minimum number of edges that one needs to add to ensure that there exist at least d followers in \mathcal{H} whose degrees are at least the degree of every leader. Then we have $b_{ALG} \leq (2 - \varepsilon)b_{OPT}$. We output that the graph \mathcal{G} contains a k -clique if $b_{ALG} \leq (2 - \varepsilon)\binom{k}{2}$. Otherwise, we output that the density of any subgraph of \mathcal{G} on k vertices is at most $\varepsilon/2$. We now prove correctness of our algorithm. We first observe that the degree of every leader in \mathcal{H} is $n + k - 1$, the degree of $a_v \in \mathcal{V}[\mathcal{H}]$ for every $v \in \mathcal{V}[\mathcal{G}]$ is n , and the degree of every other vertex is at most 1.

- (i) **Completeness:** Let $\mathcal{W} \subseteq \mathcal{V}[\mathcal{G}]$ with $|\mathcal{W}| = k$ be a clique in \mathcal{G} . Let us consider the subset $\mathcal{X} = \{a_v : v \in \mathcal{W}\} \subseteq \mathcal{V}[\mathcal{H}] \setminus \mathcal{L}$. By construction, \mathcal{X} forms an independent set in \mathcal{H} and the degree of every vertex in \mathcal{H} is n . Since, adding all the edges in $\{\{a_u, a_v\} : u, v \in \mathcal{W}, u \neq v\}$ in \mathcal{H} makes the degree of every vertex in \mathcal{X} in the resulting graph $n+k-1$, we have $b_{OPT} \leq \binom{k}{2}$. Hence, we have $b_{ALG} \leq (2 - \varepsilon)b_{OPT} \leq (2 - \varepsilon)\binom{k}{2}$.
- (ii) **Soundness:** In this case, the density of any subgraph of \mathcal{G} on k vertices is at most $\varepsilon/2$. Let $\mathcal{Y} \subseteq \mathcal{V}[\mathcal{H}] \setminus \mathcal{L}$ with $|\mathcal{Y}| = k$ be a set of any k followers. By the construction of \mathcal{H} , we have $|\mathcal{E}[\mathcal{Y}[\mathcal{H}]]| \geq (1 - (\varepsilon/2))\binom{k}{2}$. Hence, the minimum number of edges one needs to add to make the degree of every vertex in \mathcal{Y} at least $n+k-1$ is at least $k(k-1) - (\varepsilon/2)\binom{k}{2} = (2 - (\varepsilon/2))\binom{k}{2}$. In particular, we have $b_{ALG} \geq b_{OPT} \geq (2 - (\varepsilon/2))\binom{k}{2} > (2 - \varepsilon)\binom{k}{2}$.

This concludes the proof of the statement. \square

4 RESULTS FOR CORE CENTRALITY

We present our results for the HIDING LEADER problem for the core centrality measure in this section. Unlike in degree centrality case, the problem becomes NP-complete even when the core centrality of every leader is only 3. This is almost tight as we prove that the problem is polynomial time solvable if the core centrality of every leader is at most 1.

In Theorem 4.1 below, we prove that the HIDING LEADER problem is NP-complete even when the core centrality of every leader is 3. We reduce the SET COVER problem to the HIDING LEADER problem there. In the SET COVER problem, the input is a universe $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$, a collection $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$ of subsets of \mathcal{U} , and an integer t and we need to compute if there exist at most t sets

in \mathcal{S} , union of which results in \mathcal{U} . It is well known that the SET COVER problem is NP-complete [32].

THEOREM 4.1. *The HIDING LEADER problem for the core centrality measure is NP-complete even when the core centrality of every leader is 3.*

PROOF. The HIDING LEADER problem for the core centrality measure is clearly in NP. Note that computing core centrality of a node takes polynomial time [12]. To prove NP-hardness, we reduce from the SET COVER problem. Let $(\mathcal{U} = \{u_1, u_2, \dots, u_n\}, \mathcal{S} = \{S_1, S_2, \dots, S_m\}, t)$ be an instance of the SET COVER problem. To define a corresponding HIDING LEADER problem instance, we construct the graph \mathcal{G} as follows.

Intuitively, for each subset $S_i \in \mathcal{S}$, we create a path of n vertices $X_{i,1}, X_{i,2}, \dots, X_{i,n}$ in \mathcal{G} ; $(X_{i,2}, X_{i,3}), \dots, (X_{i,n-1}, X_{i,n}), (X_{i,n}, X_{i,1})$ are the edges of the above path. We also add 5 vertices $W_{i,1}$ to $W_{i,5}$ with eight edges where the four vertices in $\{W_{i,\ell} : 2 \leq \ell \leq 5\}$ form a clique with six edges; the other two edges are $(W_{i,1}, W_{i,2})$ and $(W_{i,1}, W_{i,5})$. For each $u_j \in \mathcal{U}$, we add a set of 5 vertices $\{Z_{j,\ell} : 1 \leq \ell \leq 5\}$ with eight edges where the four vertices (leaders) $\{Z_{j,\ell} : 2 \leq \ell \leq 5\}$ form a clique with six edges; the other two edges are $(Z_{j,1}, Z_{j,2})$ and $(Z_{j,1}, Z_{j,5})$. We also have an edge $(X_{i,j}, Z_{j,1})$ for every $u_j \in S_i$. We allow to add t new edges and demand that the core centrality of at least $4m + n(t+1) + t$ followers should be at least as high as the core centrality of every leader. Figure 1 illustrates the structure of our construction for sets $S_1 = \{u_1, u_2\}, S_2 = \{u_2\}, S_3 = \{u_3, u_4\}$. We now formally describe our HIDING LEADER instance.

$$\begin{aligned} \mathcal{V}[\mathcal{G}] &= \{X_{i,j} : S_i \in \mathcal{S}, u_j \in \mathcal{U}\} \cup V_1 \\ \mathcal{E}[\mathcal{G}] &= E_1 \cup E_2 \cup E_3 \cup E_4 \cup E_5 \cup E_6 \cup E_7 \cup E_8 \\ V_1 &= \{W_{i,p} : S_i \in \mathcal{S}, p \in [5]\} \cup \{Z_{j,p} : u_j \in \mathcal{U}, p \in [5]\} \\ E_1 &= \{(X_{i,j}, X_{i,j+1}) : j \in [n-1] \setminus \{1\}, i \in [m]\} \\ E_2 &= \{(X_{i,n}, X_{i,1}) : i \in [m]\} \\ E_3 &= \{(W_{i,1}, W_{i,p}) : i \in [m], p = 2, 5\} \\ E_4 &= \{(W_{i,1}, X_{i,j}) : u_j \notin S_i, i \in [m], j \in [n]\} \\ E_5 &= \{(Z_{j,1}, Z_{j,p}) : j \in [n], p = 2, 5\} \\ E_6 &= \{(X_{i,j}, Z_{j,1}) : u_j \in S_i, i \in [m], j \in [n]\} \\ E_7 &= \{(W_{i,p}, W_{i,q}) : i \in [m], 2 \leq p < q \leq 5\} \\ E_8 &= \{(Z_{j,p}, Z_{j,q}) : j \in [n], 2 \leq p < q \leq 5\} \\ \mathcal{L} &= \{Z_{j,p} : j \in [n], p = 2, 3, 4, 5\} \\ b &= t, d = 4m + n(t+1) + t \end{aligned}$$

We now claim that these two instances are equivalent. In one direction, let us assume that the SET COVER instance is a YES instance. By renaming, let us assume that the collection $\{S_1, \dots, S_t\}$ forms a valid set cover of the instance. We add the edges in the set $\mathcal{E}' = \{(X_{i,1}, X_{i,2}) : i \in [t]\}$ in the graph \mathcal{G} . Let the resulting graph be \mathcal{H} . We claim that the core centrality of every vertex in $\{Z_{i,1} : i \in [n]\} \cup \{X_{i,j} : i \in [t], j \in [n]\} \cup \{W_{i,1} : i \in [t]\} \cup \{W_{i,j} : i \in [m], j \in [4]\}$ is 3 in \mathcal{H} . We first observe that the core centrality of every leader remains 3 even after adding the edges in \mathcal{E}' . Also, for any $i \in [m]$, if the edge $(X_{i,1}, X_{i,2})$ is added in the graph, the core centrality of the $n+1$ vertices $X_{i,t}, t \in [n]$ and $W_{i,1}$ become 3. Hence after addition of the edges in \mathcal{E}' in \mathcal{G} , the core centrality of every vertex in $\{X_{i,j} : i \in [t], j \in [n]\} \cup \{W_{i,1} : i \in [t]\}$ becomes




Figure 1: Example construction for hardness from Set Cover where $\mathcal{U} = \{u_1, u_2, u_3, u_4\}$, $S = \{S_1, S_2, S_3\}$, $S_1 = \{u_1, u_2\}$, $S_2 = \{u_2\}$, $S_3 = \{u_3, u_4\}$. The red nodes are the leaders and the blue nodes are the followers.

3. Since $\{S_1, \dots, S_t\}$ forms a set cover for \mathcal{U} , the core centrality of every vertex in $\{Z_{i,1} : i \in [n]\}$ becomes 3. Lastly, the core centrality of every vertex in $\{W_{i,j} : i \in [m], j \in [4]\}$ was already 3 in \mathcal{G} and since addition of edges never decreases the core centrality of any vertex, the core centrality of these vertices are at least 3 in \mathcal{H} . Hence the HIDING LEADER instance is a YES instance.

For the other direction, let us assume that there exists a set \mathcal{E}' of edges such that in the graph $\mathcal{H} = (\mathcal{V}[\mathcal{G}], \mathcal{E}[\mathcal{G}] \cup \mathcal{E}')$, the core centrality of at least d vertices in $\mathcal{V}[\mathcal{G}] \setminus \mathcal{L}$ is at least 3; let the set of followers with core centrality at least 3 in \mathcal{H} be $\mathcal{Y} \subseteq \mathcal{V}[\mathcal{G}] \setminus \mathcal{L}$. Since adding edges in the graph never decreases the core centrality of any vertex, we have $\{W_{i,j} : i \in [m], j \in [4]\} \subseteq \mathcal{Y}$. Let us consider the following subset $\mathcal{J} \subseteq [m]$ defined as: $\mathcal{J} = \{j \in [m] : \exists 1 \leq i < k \leq n \text{ with } (X_{j,i}, X_{j,k}) \in \mathcal{F}\}$. Since $|\mathcal{F}| \leq b$ and $b = t$, we have $|\mathcal{J}| \leq t$. We claim that $\{S_j : j \in \mathcal{J}\}$ forms a set cover for \mathcal{U} . Suppose not, then at most $n-1$ vertices in $\{Z_{i,1} : i \in [n]\}$ can belong to \mathcal{Y} since $Z_{\ell,1}$ does not belong to \mathcal{Y} if u_ℓ is uncovered. Also, any vertex in $\{X_{i,\ell}, W_{i,1} : i \in [m] \setminus \mathcal{J}, \ell \in [n]\}$ does not belong to \mathcal{Y} . Hence, we have $|\mathcal{Y}| \leq 4m + t(n+1) + n-1 < d$ which contradicts our assumption that \mathcal{F} forms a valid solution for the HIDING LEADER instance. Hence the SET COVER instance is a YES instance. \square

Theorem 4.1 along with well known inapproximability result for the SET COVER problem immediately give us the following result.

COROLLARY 4.2. *There does not exist any polynomial time algorithm for approximating the number of edges one needs to add in the HIDING LEADER problem for core centrality within an approximation ratio of $(1-\alpha) \ln n$ for any constant α assuming $P \neq NP$ even when the core centrality of every leader is 3.*

PROOF. The result follows from the observation that the reduction in the proof of Theorem 4.1 is approximation preserving and the $(1-\alpha) \ln n$ inapproximability result for the SET COVER problem for any constant α assuming $P \neq NP$ [54]. \square

We now show that the hardness result in Theorem 4.1 is almost tight in the sense that if the core centrality of every leader is at

most 1 in the network, then the corresponding HIDING LEADER problem is polynomial time solvable.

PROPOSITION 4.3. *There exists a polynomial time algorithm for the HIDING LEADER problem for core centrality if the core centrality of every leader in the network is at most 1.*

PROOF. We observe that if the degree of any vertex x is at least 1, then its core centrality is at least 1. Let $\mathcal{F}' \subseteq \mathcal{F}$ be the subset of followers whose core centrality is at least 1; say $|\mathcal{F}'| = d'$. Hence the degree of every vertex in $\mathcal{F} \setminus \mathcal{F}'$ is 0. We add $\lceil (d-d')/2 \rceil$ new edges such that the degree of at least $d-d'$ vertices in $\mathcal{F} \setminus \mathcal{F}'$ becomes at least 1 in the resulting graph. We output YES if $\lceil (d-d')/2 \rceil \geq b$; otherwise we output NO. Since any optimal solution must add at least $\lceil (d-d')/2 \rceil$ edges, our algorithm is correct. \square

5 CAPTAIN NETWORKS

In this section, we show the “captain network”, originally proposed by Wanik et al. [72], also ensures that the core centrality of any leader is at most the core centrality of any captain. They propose two constructions; one for single leader and another for multiple leaders.

5.1 For Multiple Leaders

We first describe the construction in [72]. The set \mathcal{L} of leaders forms a clique. Each leader $l_i \in \mathcal{L}$ has a corresponding group of p captains $C_i = \{C_{i,1}, C_{i,2}, \dots, C_{i,p}\}$ and l_i is connected to all vertices in C_i . Assuming that $|\mathcal{L}| = h \geq 2$, there are h such sets of captains $\{C_1, C_2, \dots, C_h\}$. All vertices in the captain sets are connected as a complete h -partite graph. A captain $C_{i,j}$ serves two things: 1) It helps to hide the leader by being higher or of same centrality than the leader with maximum centrality. 2) It spreads the influence from the leader to the rest of the network. The remaining vertices $\mathcal{X} = \{X_1, X_2, \dots, X_m\}$ are each connected to one captain from each group C_i . The follower set in the network is $\mathcal{F} = \mathcal{X} \cup C_1 \cup C_2 \cup \dots \cup C_h$. Let us call the resulting graph \mathcal{G}_M . We now show that the core centrality of every leader in \mathcal{G}_M is at most the core centrality of every captain.

THEOREM 5.1. *Given a captain network \mathcal{G}_M , let $r = \lfloor \frac{m}{p} \rfloor$ denotes the minimum number of connections that a captain $C_{i,j}$ has with vertices from \mathcal{X} . Assuming we have at least 2 leaders and $p > 1$, the core centralities of the captains are either greater or same as the leaders.*

PROOF. In \mathcal{G}_M , the vertices in \mathcal{X} do not contribute in the core centrality of either the leaders or the captains. We observe that the degree of any vertex in \mathcal{X} is h . So their core centrality can be at most h . We claim that the captains and the leaders are in higher core than h . Consider an induced subgraph $G' \subset \mathcal{G}_M$ that includes only the leaders, the captains and the edges between them. In G' , the degree of any captain $C_{i,j}$ is $p(h-1) + 1$; on the other hand, the degree of any leader l_i is $h-1+p$. So, in G' , all the captains and the leaders are at least in $d_{min} = \min\{d(G, c_{i,j}), d(G, l_i)\}$ -core. This comes from the fact that all the nodes (captains and leaders) have a minimum degree of d_{min} and thus they are at least in d_{min} -core. Note that, $d(G, c_{i,j}) - d(G, l_i) = p(h-1) + 1 - (h-1+p) = (h-2)(p-1) \geq 0$

as $h \geq 2$ and $p > 1$. This implies the captains have higher degree than the leaders in G' . So, the captains have at least the same core-centrality as the leaders. Our claim is proved. Additionally, any vertex in X is in h -core and $h < d_{min}$ assuming $p > 1$. \square

COROLLARY 5.2. *Given the same captain network G , assuming $h > 2$ and $p > 1$, the core centrality of all the captains is strictly larger than the leaders.*

PROOF. The key idea is that the captains will form a core only among themselves and that will be higher core than the leaders. Now, for any captain $c_{i,j}$ the degree among themselves is $p(h - 1)$. Now that, $p(h - 1) - (h - 1 + p) = (h - 2)(p - 1) - 1 > 0$ or $(h - 2)(p - 1) > 1$ is possible when $h > 2$ and $p > 1$. So, the captains have larger core-centrality than the leaders. \square

5.2 For Single Leader

We start with the construction in [72] when $h = 1$ and show the core centralities of the leaders and the captains remain same in this case.

A single leader l ($h = 1$) has a corresponding two sets of p captains $C_1 = \{C_{1,1}, C_{1,2}, \dots, C_{1,p}\}$ and similarly it has C_2 . All vertices in the captain sets are connected as a complete bipartite graph. Each of the remaining vertices in $X = \{x_1, x_2, \dots, x_m\}$ is connected to one captain from each group C_1 and C_2 . The follower set in the network is $F = \{X \cup C_1 \cup C_2\}$.

COROLLARY 5.3. *Given the captain network described above, let $r = \lfloor \frac{m}{p} \rfloor$ denote the minimal number of connections that a captain, $c_{i,j}$ has with vertices from X . Assuming $h = 1$, the core centralities of all the captains are same as the leader.*

PROOF. The proof follows from that of Theorem 5.1. The leader has degree $2p$ where as the captains have degree $1 + p + r$. But the vertices in X has only degree 2. So the leader and the captains will be in the higher core and it will be $\min\{2p, p + 1\}$. Assuming $p \geq 2$ all the captain vertices and the leader will be in $p + 1$ -core. If $p = 1$, all the vertices in the network will be in 2-core. \square

6 SIMULATION RESULTS

In this section, we evaluate the performance of our 2 approximation algorithm in Theorem 3.2 using synthetic networks. For brevity, let us call our algorithm in Theorem 3.2 as HLDA and called the lower bound used in Theorem 3.2 as LB. We also show how well the leaders can be hidden in the captain network via the core centrality measure. Solutions were implemented in Java and experiments conducted on 3.30 GHz Intel cores with 30 GB RAM.

6.1 Evaluation of 2-Approximation Algorithm in Theorem 3.2

Settings: We generate synthetic network structures from two well-studied models: (a) Barabasi-Albert (BA) [6] and (b) Watts-Strogatz (WS) [74]. While both have “small-world” property, WS do not have a scale-free degree distribution. We generate both the datasets of 70 thousands vertices for three different edge densities: average degree of vertices as 2, 4 and 10. In the experiments we choose 20 leaders ($|\mathcal{L}| = h = 20$) randomly from the top 100 high degree vertices.

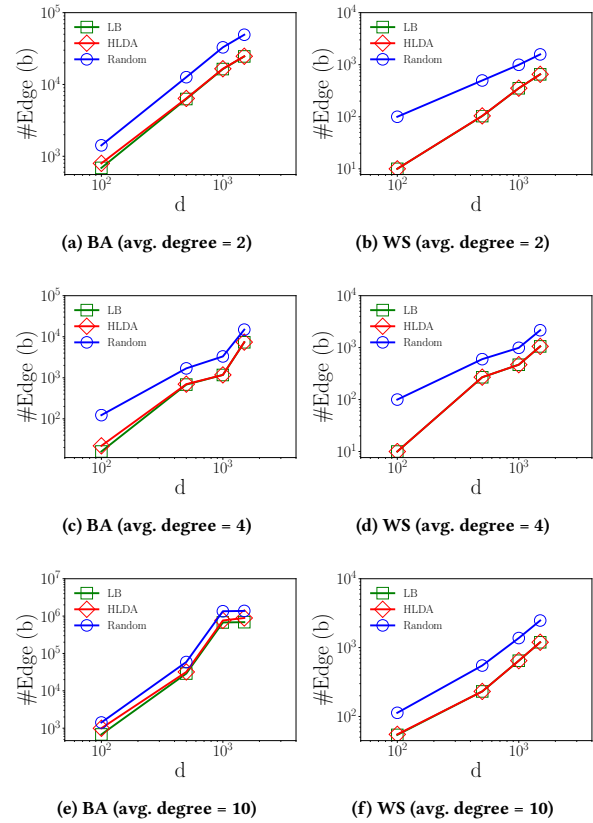


Figure 2: Number of edges added (b) by different algorithms: LB implies a loose lower bound, HLDA is our algorithm that gives 2-approximation and Random denotes a random edge addition algorithm. Clearly, in both networks while varying edge density (average degree of nodes), the number of edge addition by our algorithm HLDA is almost same as that of LB.

Baselines: We compare our algorithm (HLDA) with two baselines. Our first baseline is the lower bound used in Theorem 3.2 which we call LB. Our second baseline is Random which denotes the number of random edges one needs to add to achieve the goal. The performance metric of the algorithms is the number of edges being added to satisfy the degree centrality requirement for d followers. Hence, the quality is better when the number of edges is lower.

Results: Theorem 3.2 shows that our algorithm (HLDA) proposed for degree centrality gives a 2-approximation. However in practice it gives near optimal results. Figure 2 shows the results varying d on four datasets. Note that, the axes are in logarithmic scale. In all six datasets, the number of solution edges of HLDA is similar to LB. However, Random cannot produce high quality results. Comparing the datasets (BA and WS), the algorithms (HLDA and LB) need higher number of edges in BA as the chosen leaders (randomly chosen from 100 top degree nodes) have much higher degree than the followers due to the scale-free degree distribution.




Figure 3: Summary of the difference in core centralities between a leader and a captain in a given captain network (with 550 vertices) by varying number of captains in each group (p) and leaders (h).

6.2 Captain Networks and Core Centrality

In section 5, we prove that the core centrality of the leaders can be hidden by the captains in the captain networks [72] (Theorem 5.1 and Corollary 5.2). We empirically evaluate the core centralities of the the leaders and the captains by varying two parameters: the number of captains (p) in each group C_i and the number of leaders (h) for network with multiple leaders.

Figure 3 presents the results for a captain network with 550 vertices. For every pair of two parameters (p and h), we compute the maximum difference in core centrality between any leader and any captain. The intensity of the color signifies that the difference is higher. Higher difference also implies higher disguise for the leaders. Low values of p result into lower disguise for a leader. On the other hand, a high value of p (large number of captains in each group) with high values of h produces the maximum amount of disguise. But for a high value of p , if the number of leaders are low, i.e., low h , the amount of disguise for the leaders decreases.

We summarize our experimental findings as follows.

- ▷ HLDA produces near optimal results in practice, where as, Random cannot produce high quality results. HLDA and LB need more edges to satisfy the degree requirements for d followers in BA due to the scale-free degree distribution.
- ▷ A captain network with small number of leaders (low h) and a large number of captains in each group (high value of p) produces the maximum amount of disguise.
- ▷ A low value of p , i.e., a small number of captains in each group yields lower disguise for core centrality which is not true for other centralities such as degree, closeness, and betweenness [72].

7 CONCLUSION AND FUTURE WORK

We have shown that the HIDING LEADER problem for the core centrality measure is NP-hard to approximate with a factor of $(1 - \alpha) \ln n$ for any constant $\alpha > 0$ for optimizing the number of edges one needs to add even when the core centrality of every leader is only 3. On the other hand, we prove that the HIDING LEADER leader problem for degree centrality is polynomial time solvable if the degree of every leader is $O(1)$. Moreover, we also provide a 2 factor polynomial time approximation algorithm for the HIDING LEADER problem for optimizing the number of edges one needs to add to hide all the leaders. Hence, our results prove that, although classical complexity theoretic framework fails to compare relative difficulty of hiding leaders with respect to various centrality measures [72], hiding leaders may be significantly harder for the core centrality than the degree centrality. We complement our 2 factor approximation algorithm for the HIDING LEADER problem for degree centrality by proving that if there exists a $(2 - \epsilon)$ factor approximation algorithm for the HIDING LEADER problem for degree centrality for any constant $0 < \epsilon < 1$, then there exists a $\epsilon/2$ factor approximation algorithm for the famous Densest k -Subgraph problem which would be considered a major break through. The current best polynomial time algorithm for the Densest k -Subgraph problem achieves an approximation ratio of only $\tilde{O}(n^{1/4})$ [11]. We have also empirically evaluated our approximation algorithm which shows that our algorithm produces an optimal solution for most of the cases. We have also shown that the captain networks proposed in [72] can hide the leaders with respect to core centrality.

An important future direction is to explore the average case computational complexity of the HIDING LEADER problem for popular network centrality measures. Since the results of Waniek et al. [72] and ours establish that the HIDING LEADER problem is intractable only in the worst case, it could very well be possible that there exist heuristics that efficiently solve most randomly generated instances. If this is true, then the apparent complexity shield against manipulating various centrality measures will become substantially weak. Another immediate future work is to resolve the computational complexity of the HIDING LEADER problem for the core centrality measure when the core centrality of every leader is at most 2.

REFERENCES

- [1] Eitan Y Alimi, Lorenzo Bosi, and Chares Demetriou. 2015. *The dynamics of radicalization: a relational and comparative perspective*. Oxford University Press.
- [2] Yaniv Altshuler, Nadav Aharoni, Yuval Elovici, Alex Pentland, and Manuel Cebrian. 2013. Stealing reality: when criminals become data scientists (or vice versa). In *Security and Privacy in Social Networks*. Springer, 133–151.
- [3] Jac M Anthonisse. 1971. The rush in a graph. *Amsterdam: Mathematische Centrum* (1971).
- [4] Scott Atran, Robert Axelrod, Richard Davis, and Baruch Fischhoff. 2017. Challenges in researching terrorism from the field. *Science* 355, 6323 (2017), 352–354.
- [5] Wayne E Baker and Robert R Faulkner. 1993. The social organization of conspiracy: Illegal networks in the heavy electrical equipment industry. *Am. Sociol. Rev* (1993), 837–860.
- [6] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.
- [7] Frank M Bass. 1969. A new product growth for model consumer durables. *Manag. Sci.* 15, 5 (1969), 215–227.
- [8] Alex Bavelas. 1948. A mathematical model for group structures. *Applied anthropology* 7, 3 (1948), 16–30.
- [9] Murray A Beauchamp. 1965. An improved index of centrality. *Behavioral science* 10, 2 (1965), 161–163.
- [10] Roberta Belli, Joshua D Freilich, Steven M Chermak, and Katharine A Boyd. 2015. Exploring the crime–terror nexus in the United States: a social network analysis

- of a Hezbollah network involved in trade diversion. *Dynamics of Asymmetric Conflict* 8, 3 (2015), 263–281.
- [11] Aditya Bhaskara, Moses Charikar, Aravindan Vijayaraghavan, Venkatesan Guruswami, and Yuan Zhou. 2012. Polynomial integrality gaps for strong SDP relaxations of Densest k -subgraph. In *Proc. 23-rd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 388–405.
- [12] Kshipra Bhawalkar, Jon Kleinberg, Kevin Lewi, Tim Roughgarden, and Aneesh Sharma. 2015. Preventing unraveling in social networks: the anchored k -core problem. *SIAM J. Discrete Math.* 29, 3 (2015), 1452–1475.
- [13] Ulrik Brandes. 2001. A faster algorithm for betweenness centrality. *Journal of mathematical sociology* (2001), 163–177.
- [14] Mark Braverman, Young Kun-Ko, Aviad Rubinfeld, and Omri Weinstein. 2017. ETH Hardness for Densest- k -Subgraph with Perfect Completeness. In *Proc. 28-th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1326–1341.
- [15] David Calvey. 2017. *Covert research: The art, politics and ethics of undercover fieldwork*. Sage.
- [16] Peter J Carrington, John Scott, and Stanley Wasserman. 2005. *Models and methods in social network analysis*. Vol. 28. Cambridge university press.
- [17] Hsinchun Chen, Homa Atabakhsh, Jennifer Jie Xu, Alan Gang Wang, Byron Marshall, Siddharth Kaza, Lu Chunju Tseng, Shauna Eggers, Hemanth Gowda, Tim Petersen, et al. 2005. Coplink center: social network analysis and identity deception detection for law enforcement and homeland security intelligence and security informatics: a crime data mining approach to developing border safe research. In *Proc. 2005 National Conference on Digital government research*. Digital Government Society of North America, 112–113.
- [18] Pierluigi Crescenzi, Gianlorenzo D'Angelo, Lorenzo Severini, and Yllka Velaj. 2015. Greedily Improving Our Own Centrality in A Network. In *Proc. 14th Symposium on Experimental Algorithms*. 43–55.
- [19] Nick Crossley, Gemma Edwards, Ellen Harries, and Rachel Stevenson. 2012. Covert social movement networks and the secrecy-efficiency trade off: The case of the UK suffragettes (1906–1914). *Soc. Networks* 34, 4 (2012), 634–644.
- [20] Peter Csermely, András London, Ling-Yun Wu, and Brian Uzzi. 2013. Structure and dynamics of core/periphery networks. *Journal of Complex Networks* 1, 2 (2013), 93–123.
- [21] Gianlorenzo D'Angelo, Lorenzo Severini, and Yllka Velaj. 2016. On the Maximum Betweenness Improvement Problem. *Electronic Notes in TCS* 322 (2016), 153–168.
- [22] E. D. Demaine and M. Zadimoghaddam. 2010. Minimizing the diameter of a network using shortcut edges. In *SWAT, ser.Lecture Notes in Computer Science, H. Kaplan, Ed.* (2010), 420–431.
- [23] Fatih Demiroz and Naim Kapucu. 2012. Anatomy of a dark network: the case of the Turkish Ergenekon terrorist organization. *Trends in organized crime* 15, 4 (2012), 271–295.
- [24] Palash Dey, Pravesh K. Kothari, and Swaprava Nath. 2019. The Social Network Effect on Surprise in Elections. In *Proc. ACM India Joint International Conference on Data Science and Management of Data, 24th COMAD, 6th CODS, Kolkata, India, January 3-5, 2019*, 1–9. <https://doi.org/10.1145/3297001.3297002>
- [25] Bistra Dilkina, Katherine J. Lai, and Carla P. Gomes. 2011. Upgrading shortest paths in networks. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 76–91.
- [26] Paul AC Duijn, Victor Kashirin, and Peter MA Sloot. 2014. The relative ineffectiveness of criminal network disruption. *Scientific reports* 4 (2014), 4238.
- [27] HA Eiselt. 2018. Destabilization of terrorist networks. *Chaos, Solitons & Fractals* 108 (2018), 111–118.
- [28] Walter Enders and Paan Jindapon. 2010. Network externalities and the structure of terror networks. *J. Confl. Resolut.* 54, 2 (2010), 262–280.
- [29] Walter Enders and Xuejuan Su. 2007. Rational terrorists and optimal network structure. *J. Confl. Resolut.* 51, 1 (2007), 33–57.
- [30] Ejaz Farooq, Shoab A Khan, and Wasi Haider Butt. 2017. Covert network analysis to detect key players using correlation and social network analysis. In *Proc. 2nd International Conference on Internet of things and Cloud Computing*. ACM, 94:1–94:6.
- [31] Linton C Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.
- [32] Michael R Garey and David S Johnson. 1979. *Computers and Intractability*. Vol. 174. freeman New York.
- [33] Jacob Goldenberg, Barak Libai, and Eitan Muller. 2001. Using complex systems analysis to advance marketing theory development: Modeling heterogeneity effects on new product growth through stochastic cellular automata. *J. Acad. Mark. Sci.* 9, 3 (2001), 1–18.
- [34] Vatche Ishakian, Dóra Erdos, Evimaria Terzi, and Azer Bestavros. 2012. A Framework for the Evaluation and Management of Network Centrality. In *Proc. SIAM International Conference on Data Mining*. 427–438.
- [35] RHP Janssen and Herman Monsuur. 2012. Stable network topologies using the notion of covering. *Eur J Oper Res.* 218, 3 (2012), 755–763.
- [36] Neil F Johnson, M Zheng, Yulia Vorobyeva, Andrew Gabriel, Hong Qi, Nicolás Velásquez, Pedro Manrique, Daniela Johnson, E Restrepo, C Song, et al. 2016. New online ecology of adversarial aggregates: ISIS and beyond. *Science* 352, 6292 (2016), 1459–1463.
- [37] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proc. 9th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 137–146.
- [38] Joshua Kilberg. 2012. A basic model explaining terrorist group organizational structure. *Studies in Conflict & Terrorism* 35, 11 (2012), 810–830.
- [39] Maksim Kitsak, Lazaros K Gallos, Shlomo Havlin, Fredrik Liljeros, Lev Muchnik, H Eugene Stanley, and Hernán A Makse. 2010. Identification of influential spreaders in complex networks. *Nature physics* 6, 11 (2010), 888.
- [40] David Knoke. 2015. Emerging trends in social network analysis of terrorism and counterterrorism. *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource* (2015), 1–15.
- [41] Valdis E Krebs. 2002. Mapping networks of terrorist cells. *Connections* 24, 3 (2002), 43–52.
- [42] Yimin Lin and Kyriakos Mouratidis. 2015. Best upgrade plans for single and multiple source-destination pairs. *Geoinformatica* 19, 2 (2015), 365–404.
- [43] Roy Lindelauf, Peter Borm, and Herbert Hamers. 2009. The influence of secrecy on the communication structure of covert networks. *Soc. Networks* 31, 2 (2009), 126–137.
- [44] Kun Liu and Evimaria Terzi. 2008. Towards identity anonymization on graphs. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. ACM, 93–106.
- [45] Yong Lu, Xin Luo, Michael Polgar, and Yuanyuan Cao. 2010. Social network analysis of a criminal hacker community. *J. Comp. Inf. Sys* 51, 2 (2010), 31–41.
- [46] Ahmad Mahmood, E Charalamos, and Eli Upfal. 2016. Scalable Betweenness Centrality Maximization via Sampling. In *22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 1765–1773.
- [47] Nigel Meade and Towhidul Islam. 2006. Modelling and forecasting the diffusion of innovation—A 25-year review. *Int. J. Forecast* 22, 3 (2006), 519–545.
- [48] Sourav Medya, Petko Bogdanov, and Ambuj Singh. 2018. Making a Small World Smaller: Path Optimization in Networks. *IEEE Transactions on Knowledge and Data Engineering* 30, 8 (2018), 1533–1546.
- [49] Sourav Medya, Arlei Silva, Ambuj Singh, Prithwish Basu, and Ananthram Swami. 2018. Group centrality maximization via network design. In *Proc. 24th SIAM International Conference on Data Mining*. SIAM, 126–134.
- [50] Sourav Medya, Jithin Vachery, Sayan Ranu, and Ambuj Singh. 2018. Noticeable network delay minimization via node upgrades. *Proceedings of the VLDB Endowment* 11, 9 (2018), 988–1001.
- [51] Bisharat Rasool Memon. 2012. Identifying important nodes in weighted covert networks using generalized centrality measures. In *Proc. European Intelligence and Security Informatics Conference*. IEEE, 131–140.
- [52] Adam Meyerson and Brian Tagiku. 2009. Minimizing average shortest path distances via shortcut edge addition. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX-RANDOM)*. Springer, 272–285.
- [53] Carlo Morselli, Cynthia Giguère, and Katia Petit. 2007. The efficiency/security trade-off in criminal networks. *Soc. Networks* 29, 1 (2007), 143–153.
- [54] Dana Moshkovitz. 2015. The Projection Games Conjecture and the NP-Hardness of In n -Approximating Set-Cover. *Theory Comput.* 11 (2015), 221–235.
- [55] Evelien Otte and Ronald Rousseau. 2002. Social network analysis: a powerful strategy, also for the information sciences. *J. Inf. Sci.* 28, 6 (2002), 441–453.
- [56] D. Paik and S. Sahni. 1995. Network upgrading problems. *Networks* (1995), 45–58.
- [57] Manos Papagelis, Francesco Bonchi, and Aristides Gionis. 2011. Suggesting Ghost Edges for a Smaller World. In *International conference on Information and knowledge management (CIKM)*. 2305–2308.
- [58] N Parotisidis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2015. Selecting shortcuts for a smaller world. In *SIAM International Conference on Data Mining (SDM)*. SIAM, 28–36.
- [59] Nikos Parotisidis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2016. Centrality-Aware Link Recommendations. In *Proc. 9th International ACM Conference on Web Search and Data Mining*. 503–512.
- [60] Steve Ressler. 2006. Social network analysis as an approach to combat terrorism: Past, present, and future research. *Homeland Security Affairs* 2, 2 (2006).
- [61] Matteo Riondato and Evgenios M Kornaropoulos. 2014. Fast approximation of betweenness centrality through sampling. In *Proc. 7th International ACM Conference on Web Search and Data Mining*. 413–422.
- [62] Nancy Roberts and Sean Everton. 2016. Monitoring and Disrupting Dark Networks: A Bias toward the Center and What It Costs Us. In *Eradicating Terrorism from the Middle East*. Springer, 29–42.
- [63] Jordi Sabater and Carles Sierra. 2002. Reputation and social network analysis in multi-agent systems. In *Proc. 1st International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*. 475–482.
- [64] Marc Sageman. 2004. *Understanding terror networks*. University of Pennsylvania Press.
- [65] Stephen B Seidman. 1983. Network structure and minimum degree. *Soc. networks* 5, 3 (1983), 269–287.
- [66] Muhammad Akram Shaikh and Wang Jiabin. 2008. Network structure mining: locating and isolating core members in covert terrorist networks. *WSEAS*

- Transactions on Information Science and Applications* 5, 6 (2008), 1011–1020.
- [67] Marvin E Shaw. 1954. Group structure and the behavior of individuals in small groups. *The Journal of psychology* 38, 1 (1954), 139–149.
- [68] Malcolm K Sparrow. 1991. The application of network analysis to criminal intelligence: An assessment of the prospects. *Soc. Networks* 13, 3 (1991), 251–274.
- [69] Rachel Stevenson and Nick Crossley. 2014. Change in covert social movement networks: The “Inner Circle” of the provisional Irish Republican Army. *Soc. Mov. Stud.* 13, 1 (2014), 70–91.
- [70] Klaus Von Lampe. 2015. *Organized crime: analyzing illegal activities, criminal structures, and extra-legal governance*. Sage Publications.
- [71] Fei-Yue Wang, Kathleen M Carley, Daniel Zeng, and Wenji Mao. 2007. Social computing: From social informatics to social intelligence. *IEEE Intelligent systems* 22, 2 (2007).
- [72] Marcin Waniek, Tomasz P. Michalak, Talal Rahwan, and Michael Wooldridge. 2017. On the Construction of Covert Networks. In *Proc. 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS*. 1341–1349. <http://dl.acm.org/citation.cfm?id=3091310>
- [73] Marcin Waniek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan. 2018. Hiding individuals and communities in a social network. *Nature Human Behaviour* 2, 2 (2018), 139.
- [74] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of “small-world” networks. *Nature* 393, 6684 (1998), 440.
- [75] Jennifer Xu and Hsinchun Chen. 2005. Criminal network analysis and visualization. *Commun. ACM* 48, 6 (2005), 100–107.
- [76] Yuichi Yoshida. 2014. Almost linear-time algorithms for adaptive betweenness centrality using hypergraph sketches. In *Proc. 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1416–1425.
- [77] Elena Zheleva and Lise Getoor. 2009. To Join or Not to Join: The Illusion of Privacy in Social Networks with Mixed Public and Private User Profiles. In *Proceedings of the 18th International Conference on World Wide Web*. ACM, 531–540.