

How Members of Covert Networks Conceal the Identities of Their Leaders

MARCIN WANIEK, New York University Abu Dhabi, UAE and University of Warsaw, Poland
TOMASZ P. MICHALAK, University of Warsaw, Poland
MICHAEL WOOLDRIDGE, University of Oxford, UK
TALAL RAHWAN, New York University Abu Dhabi, UAE

Centrality measures are the most commonly advocated social network analysis tools for identifying leaders of covert organizations. While the literature has predominantly focused on studying the effectiveness of existing centrality measures or developing new ones, we study the problem from the opposite perspective, by focusing on how a group of leaders can avoid being identified by centrality measures as key members of a covert network. More specifically, we analyze the problem of choosing a set of edges to be added to a network to decrease the leaders' ranking according to three fundamental centrality measures, namely, degree, closeness, and betweenness. We prove that this problem is NP-complete for each measure. Moreover, we study how the leaders can construct a network from scratch, designed specifically to keep them hidden from centrality measures. We identify a network structure that not only guarantees to hide the leaders to a certain extent but also allows them to spread their influence across the network.

CCS Concepts: • **Mathematics of computing** → **Graph algorithms**; • **Security and privacy** → *Social aspects of security and privacy*;

Additional Key Words and Phrases: Social networks, covert networks, centrality, complexity analysis.

ACM Reference format:

Marcin Waniek, Tomasz P. Michalak, Michael Wooldridge, and Talal Rahwan. 2021. How Members of Covert Networks Conceal the Identities of Their Leaders. *ACM Trans. Intell. Syst. Technol.* 13, 1, Article 12 (November 2021), 29 pages.

<https://doi.org/10.1145/3490462>

1 INTRODUCTION

Mapping terrorist networks is of vital importance to any counter-terrorism efforts. Not only does this help to understand their operational structure and modus operandi, but it also plays a key

Marcin Waniek was supported by the Polish National Science Centre Grant No. 2015/17/N/ST6/03686. Michael Wooldridge was supported by the European Research Council under Advanced Grant No. 291528 (“RACE”). Tomasz Michalak was supported for this version of this work by the Polish National Science Centre Grant No. 2016/23/B/ST6/03599, and for the previous, conference version by the Polish National Science Centre Grant No. 2014/13/B/ST6/01807 and the European Research Council under Advanced Grant No. 291528 (“RACE”).

Authors' addresses: M. Waniek, New York University Abu Dhabi, Saadiyat Island, 129188, Abu Dhabi, UAE and University of Warsaw, Banacha 2, 02-097 Warsaw, Poland; email: mjwaniek@nyu.edu; T. P. Michalak, University of Warsaw, Banacha 2, 02-097 Warsaw, Poland; email: tpm@mimuw.edu.pl; M. Wooldridge, University of Oxford, Wolfson Building, Parks Road, OX1 3QD, Oxford, UK; email: mjw@cs.ox.ac.uk; T. Rahwan, New York University Abu Dhabi, Saadiyat Island, 129188, Abu Dhabi, UAE; email: tr72@nyu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2157-6904/2021/11-ART12 \$15.00

<https://doi.org/10.1145/3490462>

role in designing and implementing destabilization strategies [11, 24, 48]. One of the most common strategies requires identifying key individuals who are suspected to play central roles in the organization [22, 51]. This task can be completed using *centrality measures*—metrics developed in graph theory to quantify the importance of nodes in networks [29, 43, 46]. Arguably, the three fundamental such measures are: (i) *degree centrality*, which ranks each node based on the number of neighbours it has; (ii) *closeness centrality*, which ranks each node based on its average distance to other nodes; and (iii) *betweenness centrality*, which ranks each node based on the relative number of shortest paths that go through that node.

Unfortunately, understanding how criminals organize themselves in a network is challenging at various levels [40, 55]; the data may be incomplete, the nature of the relationship between two criminals may be unclear, and the network may evolve continuously. The literature on this research problem generally agrees that criminals in general, and terrorists, in particular, face a tradeoff between *secrecy* and *efficiency* [49]. Overall, two approaches in this literature can be distinguished, which we briefly discuss next.

In the first approach, researchers study known topologies of historical or contemporary criminal networks, with the aim being to understand why particular structures have emerged [14, 16, 38]. Perhaps the most comprehensive study in this body of research is due to Kilberg [38], who analyzed an extensive dataset of more than 240 terrorist networks, and provided a classification of those networks based on their structure and functionality. Furthermore, using regression analysis, the author tried to quantify the degree to which the structure of terrorist networks is influenced by such variables as the Gross Domestic Product level of the target country, the political rights and civil liberties therein, and the inclination to attack police and military targets in that country.

In this article, we contribute to the second approach in the literature, which is more theoretical in nature. In particular, the aim of this approach is to explain the structural properties of covert networks by developing explicit models of the terrorists' preferences as well as the different choices they face [20, 34, 42]. With such analysis, certain network topologies typically emerge as the result of modelling the terrorists as rational decision makers. A notable example of such a model is that of Lindelauf et al. [42], who consider the tradeoff between the secrecy and operational efficiency of a terrorist network and borrow concepts from both game theory and graph theory to identify alternative topologies. The authors argue that it is inefficient if a message has to be passed several times from one person to another, i.e., if the shortest path from the sender to the receiver is relatively long. Based on this observation, the authors define efficiency as the (normalized) reciprocal of the *total distance* of the graph, i.e., the sum of shortest distances between any two nodes in the network. However, the authors consider the secrecy of a given node to be proportional to the fraction of the network that remains unexposed when that node is detected. The secrecy of the entire network is then the sum of the secrecy scores of all the nodes therein.

In this article, we also propose a theoretical model to study the secrecy-efficiency tradeoff in covert networks. Our model differs from previous ones in several ways. First, inspired by studies of real-life covert networks [12, 44], we take a *leader-centric approach*, i.e., we focus on the role played in terrorist networks by their leaders. In more detail, we investigate how the topology of the network could be deliberately designed to keep the identity of the leader(s) hidden. Compared to the literature on identifying leaders of terrorist networks, which typically assumes that such leaders are oblivious to the techniques and methods used by law enforcement agencies, we assume that this is not the case, i.e., the terrorist leaders in our model *strategically shape their network to shield themselves from being detected by centrality measures*. In fact, recent media reports and academic studies of criminal and terrorist organizations suggest that members of such organizations are becoming increasingly tech-savvy [36, 50]. Hence, their obliviousness to the available social network analysis techniques should not be taken for granted.

As already argued, secrecy is not the only objective that the leaders of a terrorist network may have. Indeed, if they were concerned only with hiding themselves, then they would simply cut most (if not all) of their connections in the network. This, however, would clearly impair the leaders' efficiency. In our model, the efficiency of the leaders is defined as their *influence over the network*. In other words, the leaders in our model face the tradeoff between hiding from centrality measures and influencing the network members. Note that a node's influence over a network can be quantified according to various models, most notably the *independent cascade* model [27] and the *linear threshold* model [37].

In the first part of the article, we focus on the computational aspects of modifying an existing network to hide the leaders from centrality analysis. More specifically, we analyze the hardness of identifying a set of edges to be added between the followers so that the *ranking* of the leaders (based on one of the three main centrality measures) drops below a certain threshold. At first glance, this problem may seem easy at least for degree centrality, which is mathematically uninvolved. Indeed, it is straightforward to decrease the *value* of degree centrality; it simply involves removing arbitrary edges of the leader [65]. Surprisingly, however, we find that our problem of decreasing the *ranking* of a node according to degree centrality is much more challenging. More precisely, Theorem 1 states that the aforementioned problem is NP-complete for degree centrality, whereas Theorems 2 and 3 provide the same result, but for closeness and betweenness centrality, respectively.

Given this hardness of modifying an existing network, we turn our attention to a different question, which is *how a terrorist network could be built from scratch so that the leaders are hidden and, at the same time, have a reasonable influence over the network members*. Here, the main goal is to identify a network structure in which the leaders surround themselves with an "inner circle" of trustees, called "*captains*," whose role is to conceal the identity of the leaders and to pass on their influence to the rest of the network. We identify one such network structure and prove that every captain is guaranteed to be ranked higher than any of the leaders (according to the three standard centrality measures). In fact, "inner circles" have been identified in various real-life terrorist networks such as, e.g., Al-Qaeda [4] and IRA [56]. While we do not have access to data that confirms that those real-life "inner circles" have a similar structure to the ones obtained in this article, we hope that our results shed more light on why such circles may exist in covert networks.

A preliminary version of this work was published in the Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2017). The new results include: (i) the NP-completeness proof of the problem of Hiding Leaders given the betweenness centrality (Theorem 3, pages 11–14); (ii) the minimization version of the problem of Hiding Leaders, and the proof that an optimal solution to this problem cannot be approximated better than logarithmically given the closeness and betweenness centralities (Section 3.2, pages 14–16); (iii) we now show that the network topology considered in this study can be embedded into a larger structure without compromising the safety of the leaders (Section 4.2, pages 19–21); (iv) we now significantly extend the numerical analysis of the interplay between the influence and centrality of the leaders in our networks, including the analysis for additional centrality measures (Section 5.1, pages 21–23); and (v) we conduct experiments investigating the attack tolerance of captain networks in comparison to other types of network structures (Section 5.2, pages 23–26).

1.1 Related Work

Our work is closely related to the literature about manipulating centrality measures in social networks. The problem of strategically lowering centrality values in existing networks was shown to be computationally intractable for most centrality measures, although some simple heuristics

proved to be surprisingly effective [65]. The setting in which the entity analyzing the networks is aware of the hiding attempts of the network members was also considered [66]. Another work presented a set of intuitive axioms that should characterize a centrality measure that is hard to manipulate [68]. The hiding problem was also analyzed for non-standard network models, such as multilayer networks [64], in which connections of different types may exist in the same structure. Dey and Medya [17] analyzed the problem of hiding network leaders, proposed in the preliminary version of this work, for the core centrality; they also presented a deeper theoretical analysis of the approximation version of the problem given the degree centrality. While all of the works mentioned above focus on lowering the centrality of a given node in the hope of concealing it from an outside observer, the problem of strategically increasing centrality has also been considered [13] with centrality being treated as a proxy of popularity.

Evasion techniques similar to these considered in our work for centrality measures were proposed for several other social network analysis tools [35], in many cases with the goal of privacy protection. Closely cooperating groups of nodes might want to avoid identification by community detection algorithms [65], especially since such algorithms have been used to infer private information about social media users, including their sexual orientation [52]. Similarly, some people might prefer to avoid publicly disclosing certain relationships and actively prevent their detection by link prediction algorithms [67, 71, 73]. Other types of social network analysis tools for which hiding or evasion techniques were developed include source detection algorithms [63] and node similarity measures [18].

Finally, our work can be related to the field of adversarial machine learning [31, 41], particularly to so-called “evasion attacks” that mislead the machine learning models by modifying their input data. Applications of the evasion attacks range from causing autonomous vehicles to misclassify road signs [23, 57], through creating email messages that can fool spam filters [7, 10], to concealing malicious code within network packages [28]. A subarea of this literature focuses on adversarial attacks on networks [58], where network modification are used to disrupt the results of machine learning algorithm, such as algorithms for node classification [15, 62, 74] or node embedding [8, 72]. It is worth noting that the techniques used in this literature are based on exploiting flaws in the machine learning algorithms, and thus cannot be directly applied to social network analysis methods considered in this work.

2 PRELIMINARIES

In this section, we present basic notation and concepts that will be used throughout the article. For the convenience of the reader, Table 1 provides a summary of the notation used in the article.

2.1 Basic Network Notation

Let us denote by $G = (V, E) \in \mathbb{G}$ a network, where $V = \{v_1, \dots, v_n\}$ is the set of n nodes and $E \subseteq V \times V$ is the set of edges. We denote an edge between nodes v_i and v_j by (v_i, v_j) . In this article, we consider *undirected* networks, meaning that E is a set of unordered pairs, i.e., we do not discern between edges (v_i, v_j) and (v_j, v_i) . We also assume that networks do not contain self-loops, i.e., $\forall v_i \in V (v_i, v_i) \notin E$.

A path in a network $G = (V, E)$ is an ordered sequence of distinct nodes, $p = \langle v_{i_1}, \dots, v_{i_k} \rangle$, in which every two consecutive nodes are connected by an edge in E . We consider the length of a path to be the number of edges in that path. The set of all shortest paths between a pair of nodes, $v_i, v_j \in V$ will be denoted by $\Pi_G(v_i, v_j)$. The distance between a pair of nodes $v_i, v_j \in V$, i.e., the length of a shortest path between them, is denoted by $d_G(v_i, v_j)$. Furthermore, a network is said to be *connected* if and only if there exists a path between every pair of nodes in that network.

Table 1. Notation Used in the Article

Notation	Meaning
$ X $	The number of elements of set X
$V = \{v_1, \dots, v_n\}$	The set of network's nodes
E	The set of network's edges
$\Pi_G(v_i, v_j)$	The set of shortest paths between v_i and v_j in the network G
$d_G(v_i, v_j)$	The length of a shortest path between v_i and v_j in the network G
$N_G(v_i)$	The set of neighbors of v_i in the network G
$N_G(v_i, v_j)$	The set of common neighbors of v_i and v_j in the network G
$c_{dg}(G, v_i)$	The degree centrality of v_i in the network G
$c_{cl}(G, v_i)$	The closeness centrality of v_i in the network G
$c_{bt}(G, v_i)$	The betweenness centrality of v_i in the network G
$\text{inf}_G(v_i, v_j)$	The influence of v_i on v_j in the network G
$\text{inf}_G(v_i)$	The influence of v_i over the entire network G
L	The set of leaders of the network
F	The set of followers in the network
b	The maximum number of edges that can be added to the network
δ	The safety margin of the leaders
\hat{A}	The set of edges that can be added to the network
A^*	The set of edges added to the network
C	The set of captain in a captain network

We denote by $N_G(v_i)$ the set of *neighbors* of v_i in G , i.e., $N_G(v_i) = \{v_j \in V : (v_j, v_i) \in E\}$. Finally, we denote by $N_G(v_i, v_j)$ the set of common neighbors of nodes v_i and v_j , i.e., $N_G(v_i, v_j) = N_G(v_i) \cap N_G(v_j)$.

To make the notation more readable, we will often denote two arbitrary nodes by v and w , instead of v_i and v_j . Moreover, we will often omit the network itself from the notation whenever it is clear from the context, e.g., by writing $d(v, w)$ instead of $d_G(v, w)$. This applies not only to the notation presented thus far but rather to all notation in this article.

2.2 Centrality Measures

The notion of *centrality* in human organizations was first introduced by Bavelas [5]. Intuitively, a centrality measure is a function, $c : \mathbb{G} \times V \rightarrow \mathbb{R}$, that expresses the relative importance of any given node in any given network. Arguably, the three best-known centrality measures are *degree*, *closeness* and *betweenness* [26]. Due to their simple closed-form formulas, they are very amenable to theoretical analysis, and we focus our attention on them in Sections 3 and 4 of our work.

Degree centrality was introduced by Shaw [54]. It assumes that the importance of a node is proportional to the number of its neighbors. Formally, the degree centrality of a node $v_i \in V$ in a network G is defined as follows:

$$c_{dg}(G, v_i) = \frac{|N(v_i)|}{n - 1}.$$

Closeness centrality, introduced by Beauchamp [6], quantifies the importance of a node in terms of shortest distances from this node to all other nodes in the network. As such, the most important node is the one with the shortest average path length to all other nodes. The normalized closeness centrality of $v_i \in V$ in a connected network G can be expressed as

$$c_{cl}(G, v_i) = \frac{n - 1}{\sum_{v_j \in V} d(v_i, v_j)}.$$

Betweenness centrality was developed independently by Anthonisse [2] and Freeman [25]. This measure quantifies the importance of a given node in the context of network flow. In more detail, if we consider all the shortest paths in the network, then the importance of any given node increases with the number of such paths that go through that node. The normalized betweenness centrality of a node $v_i \in V$ in a connected network G can be expressed as

$$c_{bt}(G, v_i) = \frac{2}{(n-1)(n-2)} \sum_{v_j, v_k \in V \setminus \{v_i\}} \frac{|\{p \in \Pi(v_j, v_k) : v_i \in p\}|}{|\Pi(v_j, v_k)|}.$$

The remaining centrality measures are defined more intricately, and as such are less amenable to theoretical analysis. We consider them in Section 5 of our work.

Eigenvector centrality [9] quantifies the importance of a given node based on the importance of its neighbors. Formally, it is defined for a node v_i in a network G as

$$c_{eg}(G, v_i) = \chi_i^*,$$

where χ^* is the eigenvector corresponding to the largest eigenvalue of the adjacency matrix of the network.

Hyperlink-induced Topic Search centrality (HITS) [39] quantifies the importance of a given node based on interpreting the network structure as connections between web pages. Each node is then assigned a score as a hub (a page that links to many pages) and as an authority (a page that many pages link to). In our analysis, we use the authority score of a node, i.e., the HITS centrality of a node $v_i \in V$ in a network G can be expressed as

$$c_{hi}(G, v_i) = x_i,$$

where x_i is computed using the normalized iterative algorithm from Kleinberg [39].

2.3 Models of Influence

The propagation of influence through the network is often described in terms of node activation. When a certain node is sufficiently influenced by its neighbors, it becomes “active,” and starts influencing any “inactive” neighbors, and so on. To initiate this propagation process, a set of nodes (known as the *seed set*) must be activated right from the start. Assuming that time moves in discrete rounds, we denote by $I(t) \subseteq V$ the set of nodes that are active at round t , implying that $I(1)$ is the seed set. The way influence propagates to inactive nodes depends on the influence model under consideration. Arguably, the two main models of influence are:

- *Independent Cascade* [27]: In this model, every pair of nodes is assigned an activation probability, $p : V \times V \rightarrow [0, 1]$. Then, in every round, $t > 1$, every node $v_i \in V$ that became active in round $t - 1$ activates every inactive neighbor, $v_j \in N(v_i) \setminus I(t - 1)$, with probability $p(v_i, v_j)$. The process ends when there are no newly activated nodes, i.e., when $I(t) = I(t - 1)$.
- *Linear Threshold* [37]: In this model, every node $v_i \in V$ is assigned a *threshold value*, t_{v_i} , which is sampled (according to some probability distribution) from the set $\{0, \dots, |N(v_i)|\}$. Then, in every round, $t > 1$, every inactive node v_i becomes active, i.e., becomes a member of $I(t)$, if $|I(t - 1) \cap N(v_i)| \geq t_{v_i}$. The process ends when there are no newly activated nodes, i.e., when $I(t) = I(t - 1)$.

In either model, the influence of a node, v_i , on another node, v_j , is denoted by $\text{inf}_G(v_i, v_j)$ and is defined as *the probability that v_j gets activated given the seed set $\{v_i\}$* . We assume that $\text{inf}_G(v_i, v_i) = 0$ for all $v_i \in V$. The influence of v_i over the entire network G is then defined as $\text{inf}_G(v_i) = \sum_{v_j \in V} \text{inf}_G(v_i, v_j)$.

Finally, it has been proposed to approximate influence of nodes in the network using the *Shapley centrality* [45, 59]. This game-theoretic centrality quantifies the importance of a given node based on its Shapley value in a particular game defined over the network. The Shapley value [53] is a uniquely fair assignment of payoff in a coalitional game, and, for player v_i in a game with utility function u , it is defined as

$$\psi(v_i, u) = \sum_{C \subseteq V \setminus \{v_i\}} \frac{|C|!(n - |C| - 1)!}{n!} (u(C \cup \{v_i\}) - u(C)).$$

The Shapley centrality of a given node v_i in a network G is then defined as

$$c_{sh}(G, v_i) = \psi(v_i, u^*),$$

where $u^* = |C \cup \bigcup_{v_i \in C} N(v_i)|$. In words, it is the Shapley value of node v_i in the game in which the value of a coalition equals the number of nodes in this coalition and all their neighbors. The intuition behind using the Shapley centrality as a proxy for influence is that the Shapley value of the above game represents an average number of new neighbors that node v_i brings to a coalition. If v_i has a relatively high Shapley value, then v_i is more likely to have some degree of exclusivity in being connected to some nodes than many other nodes. Hence, it is v_i that is more likely to be the one that “activates” others.

3 PROBLEM STATEMENT AND THEORETICAL ANALYSIS

In this section, we state the main theoretical problem of this work and prove its NP-completeness.

As mentioned in the Introduction, we assume that the terrorist network is composed of two types of agents, namely, the *leaders* and the *followers*. Furthermore, we assume that the leaders are aware that law-enforcement agencies may use centrality analysis to identify them. Thus, the leaders would like to strategically modify the existing network so that their centrality becomes lower than a certain, predefined threshold $\delta \in \mathbb{N}$ that we refer to as a *safety margin*. To achieve this objective, no more than $b \in \mathbb{N}$ modifications can be made to the network (b can be thought of as a “budget” to spend). Since removing edges would mean that existing communication links are severed, we assume that the network can be modified only by adding edges. Furthermore, since adding an edge to any leader increases that leader’s degree centrality, we assume that edges can only be added between followers. Formally, we define *the problem of Hiding Leaders* as follows:

Definition 1 (Hiding Leaders). This problem is defined by a tuple, (G, L, b, c, δ) , where $G = (V, E) \in \mathbb{G}$ is a network, $L \subset V$ is a set of leaders to be hidden, $b \in \mathbb{N}$ is a budget specifying the maximum number of edges that can be added, $c : \mathbb{G} \times V \rightarrow \mathbb{R}$ is a centrality measure, and $\delta \in \mathbb{N}$ is a chosen safety margin. Then, if we denote by $F = V \setminus L$ the set of “followers,” then the goal is to identify a set of edges to be added to the network, $A^* \subseteq F \times F$, such that $|A^*| \leq b$ and the resulting network $G' = (V, E \cup A^*)$ contains at least δ followers that each have a centrality score higher than that of any leader, i.e.,

$$\exists F' \subseteq F |F'| \geq \delta \wedge \forall f \in F' \forall l \in L (c(G', f) > c(G', l)).$$

In other words, the goal is to identify a set of connections between the followers, the addition of which ensures that all leaders are safely hidden from the given centrality measure.

3.1 Computational Complexity Analysis

Intuitively, the above problem should be easy to solve for the degree centrality measure. Indeed, adding an edge between any two (disconnected) followers increases their degree centrality with respect to all the leaders. However, we prove below that the problem is in fact NP-complete for the degree centrality measure.

THEOREM 1. *The problem of Hiding Leaders is NP-complete given the degree centrality.*

PROOF. The problem is trivially in NP, since after the addition of a given A^* it is possible to compute the degree centrality for all nodes in polynomial time.

Next, we prove that the problem is NP-hard. To this end, we propose a reduction from the NP-complete problem of *Finding k -clique*. The decision version of this problem is defined by a network, $G = (V, E)$, and a constant, $k \in \mathbb{N}$, where the goal is to determine whether there exist k nodes in G that form a clique.

Let us assume that $k \geq 3$ (if $k = 2$ then the problem is trivial). Given an instance of the problem of *Finding k -clique* by some $k \geq 3$ and a network $G = (V, E)$, let us construct a network, $H = (V', E')$, as follows:

- **The set of nodes:** For every node, $v_i \in V$, we create a single node, v_i , as well as $|N_G(v_i)|$ other nodes, denoted by $X = \{x_{i,1}, \dots, x_{i,|N_G(v_i)|}\}$. Additionally, we create one node called y , as well as $n + k - 1$ other nodes, namely, $L' = l_1, \dots, l_{n+k-1}$;
- **The set of edges:** We create an edge between two nodes $v_i, v_j \in V$ if and only if this edge was not present in G , i.e., $(v_i, v_j) \in E' \iff (v_i, v_j) \notin E$. Additionally, for every v_i , we create an edge (v_i, y) as well as an edge $(v_i, x_{i,j})$ for every $x_{i,j}$. We also create an edge (l_i, l_j) between every pair of nodes $l_i, l_j \in L'$, except for the edge (l_1, l_2) . Finally, we create two additional edges, (l_1, y) and (l_2, y) .

An example of such an H network is illustrated in Figure 1. Now, consider the following instance of the problem of hiding leaders, (H, L, b, c, δ) , where:

- $H = (V', E')$ is the network we just constructed;
- $L = V' \setminus V$;
- $b = \frac{k(k-1)}{2}$;
- c is the degree centrality measure;
- $\delta = k$.

Next, we reduce the problem of Finding k -cliques in G to the aforementioned instance of Hiding Leaders in H . To this end, from the definition of the problem of Hiding Leaders, we know that the edges to be added to H must be chosen from $F \times F$. Since in our instance we have: $F = V' \setminus L = V' \setminus (V' \setminus V) = V$, then the edges to be added to H must be chosen from $V \times V$. However, since the edges in $(V \times V) \setminus E$ are already present in H (see how H is created), then the edges to be added to H must be chosen from E . Out of those edges, we need to choose a subset, $A^* \subseteq E$, as a solution to the problem. In what follows, we will show that a solution to the above instance of the problem of Hiding Leaders in H corresponds to a solution to the problem of *Finding k -clique* in G .

First, note that each of the k nodes with the highest degree centrality in H must be a member of L' . This is because there are more than k nodes in L' , each of which has a degree of $n + k - 2$, while the degree of every node in $V' \setminus L'$ is smaller than $n + k - 2$. Thus, for A^* to be a solution to the problem of hiding leaders, the addition of A^* to H must increase the degree of at least k nodes in V such that each of them has a degree of at least $n + k - 1$ (note that the addition of A^* only increases the degrees of nodes in V , as we already established that $A^* \subseteq E$). Now, since in H the degree of every node in V equals n (because of the way H is created), then to increase the degree of k such nodes to $n + k - 1$, each of them must be an end of *at least* $k - 1$ edges in A^* . But, since the budget in our problem instance is $\frac{k(k-1)}{2}$, then the only possible choice of A^* is the one that increases the degree of *exactly* k nodes in V by *exactly* $k - 1$. If such a choice of A^* is available, then surely those k nodes would form a clique in G , since all $\frac{k(k-1)}{2}$ edges in A^* are taken from G . \square

Having proven the NP-completeness of the problem given the degree centrality, we next prove its NP-completeness given the closeness centrality.

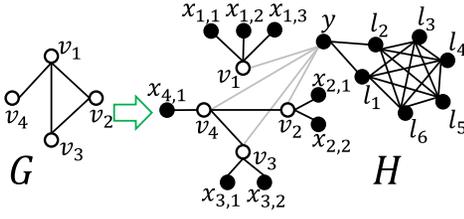


Fig. 1. An illustration of the network used in the NP-completeness proof of the problem of Hiding Leaders given the degree centrality.

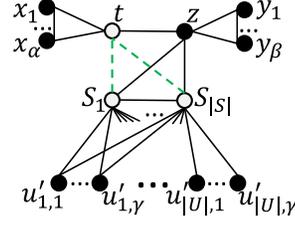


Fig. 2. An illustration of the network used in the NP-completeness proof of the problem of Hiding Leaders given the closeness and betweenness centrality.

THEOREM 2. *The problem of Hiding Leaders is NP-complete given the closeness centrality.*

PROOF. The problem is trivially in NP, since after the addition of a given A^* it is possible to compute the closeness centrality of all nodes in polynomial time.

Next, we prove that the problem is NP-hard. To this end, we propose a reduction from the NP-complete *3-Set-Cover* problem. The decision version of this problem is defined by a universe $U = \{u_1, \dots, u_{|U|}\}$ and a collection of sets $S = \{S_1, \dots, S_{|S|}\}$ such that $\forall_i S_i \subset U \wedge |S_i| = 3$, where the goal is to determine whether there exist b elements of S the union of which equals U .

Given an instance of the *3-Set-Cover* problem, let us construct a network G as follows:

- **The set of nodes:** For every $S_i \in S$, we create a single node denoted by S_i . For every $u_i \in U$, we create γ nodes denoted by $u'_{i,1}, \dots, u'_{i,\gamma}$, where $\gamma = |S| + 1$. We denote the set of all S_i nodes by S , and the set of all $u'_{i,j}$ nodes by U' . In addition, we create α nodes denoted by $X = \{x_1, \dots, x_\alpha\}$, where $\alpha = 4|S| + 4$, and create β nodes denoted by $Y = \{y_1, \dots, y_\beta\}$, where $\beta = 3|S| + 4$. Last, we create two additional nodes, t and z .
- **The set of edges:** First, we create the edge (t, z) . Then, for every node x_i we create an edge (x_i, t) , for every node y_i we create an edge (y_i, z) , and for every node $S_i \in S$, we create an edge (S_i, z) . Moreover, for every node $u'_{j,k}$, we create an edge $(S_i, u'_{j,k})$ if and only if $u_j \in S_i$. We also create edges such that the nodes in X form a clique, i.e., we create an edge (x_i, x_j) for every $x_i, x_j \in X$. Likewise, we create edges such that the nodes in Y form a clique, i.e., we create an edge (y_i, y_j) for every $y_i, y_j \in Y$. Finally, we create edges such that the nodes in S form a clique, i.e., we create an edge (S_i, S_j) for every $S_i, S_j \in S$.

An example of the resulting network, G , is illustrated in Figure 2. Now, consider the following instance of the problem of hiding leaders, (G, L, b, c, δ) , where:

- G is the network we just constructed;
- $L = \{z\} \cup X \cup Y \cup U'$;
- b is the parameter of the *3-Set-Cover* problem (where the goal is to determine whether there exist b elements of S the union of which equals U);
- c is the closeness centrality measure;
- $\delta = 1$.

From the definition of the problem of Hiding Leaders, we see that the only edges that can be added to the graph are edges between t and the members of S , i.e., $A^* \subseteq \hat{A}$, where $\hat{A} = \{(t, S_1), \dots, (t, S_k)\}$. Notice that any such a choice of A^* corresponds to selecting a subset of $|A^*|$ elements of S in the *3-Set-Cover* problem. In what follows, we will show that a solution to the above instance of Hiding Leaders corresponds to a solution to the *3-Set-Cover* problem.

Table 2. Distances between Nodes in the Graph Constructed for the Closeness Centrality Proof

v	$d(v, z)$	$d(v, t)$	$\sum_{x_i \in X} d(v, x_i)$	$\sum_{y_i \in Y} d(v, y_i)$	$\sum_{u'_{i,j} \in U'} d(v, u'_{i,j})$	$\sum_{S_i \in S} d(v, S_i)$
z	0	1	2α	β	$2 U \gamma$	$ S $
t	1	0	α	2β	$(3 U - U_{A^*})\gamma$	$2 S - A^* $
x_i	2	1	$\alpha - 1$	3β	$(4 U - U_{A^*})\gamma$	$3 S - A^* $
y_i	1	2	3α	$\beta - 1$	$3 U \gamma$	$2 S $
u_i	2	≥ 2	$\geq 3\alpha$	3β	$\geq 2(U \gamma - 1)$	$\geq S $
S_i	1	≥ 1	$\geq 2\alpha$	2β	$3\gamma + 2(U - 3)\gamma$	$ S - 1$

First, we show that for every $v \in V \setminus \{t, z\}$ and every $A^* \subseteq \hat{A}$, we either have $c(G', v) < c(G', t)$ or $c(G', v) < c(G', z)$, where $G' = (V, E \cup A^*)$. To this end, we show that the following holds, where $D(G', v) = \frac{n-1}{c(G', v)} = \sum_{w \in V \setminus \{v\}} d(v, w) : \forall v \in V \setminus \{t, z\} \forall_{A^* \subseteq \hat{A}} D(G', v) > D(G', t) \vee D(G', v) > D(G', z)$.

Let U_{A^*} denote the elements of U covered by sets from S the representations of which are connected with t by edges from A^* , i.e., $U_{A^*} = \{u_j \in U : \exists_{(t, S_i) \in A^*} u_j \in S_i\}$. Table 2 presents the distances between different nodes in G' . Based on the information from the table, we have the following:

- $D(G', z) = 2\alpha + \beta + 2|U|\gamma + |S| + 1$;
- $D(G', t) = \alpha + 2\beta + (3|U| - |U_{A^*}|)\gamma + 2|S| - |A^*| + 1$;
- $D(G', x_i) = \alpha + 3\beta + (4|U| - |U_{A^*}|)\gamma + 3|S| - |A^*| + 2 > D(G', t)$;
- $D(G', y_i) = 3\alpha + \beta + 3|U|\gamma + 2|S| + 2 > D(G', z)$;
- $D(G', u_i) \geq 3\alpha + 3\beta + 2|U|\gamma + |S| + 2 > D(G', z)$;
- $D(G', S_i) \geq 2\alpha + 2\beta - 3\gamma + 2|U|\gamma + |S| + 1 > D(G', z)$, because $\beta > 3\gamma$.

Therefore, either t or z has the highest closeness centrality. Since $z \in L$, $t \in F$ and the safety margin $\delta = 1$, then $A^* \subseteq \hat{A}$ is a solution to the constructed instance of the Hiding Leaders problem if and only if $D(G', t) < D(G', z)$. We have that: $D(G', z) - D(G', t) = \alpha - \beta + (|U_{A^*}| - |U|)\gamma + |A^*| - |S|$.

We will now prove that if there exists a solution S^* to the given instance of the 3-Set-Cover problem, then there also exists a solution to the constructed instance of the Hiding Leaders problem. To this end, let $A^* = \{(t, S_i) : S_i \in S^*\}$. We will show that A^* is a solution to the constructed instance of the Hiding Leaders problem. Since S^* is a solution to the given instance of the 3-Set-Cover problem, then $U_{A^*} = U$, as all elements of U are covered by sets in S^* . Hence, given that $|A^*| \geq 1$, we have: $D(G', z) - D(G', t) \geq \alpha - \beta + 1 - |S|$. By substituting values of α and β we get: $D(G', z) - D(G', t) \geq |S| + 1 - |S| > 0$. Therefore, we have $D(G', t) < D(G', z)$, which, as mentioned above, is a sufficient condition for A^* to be a solution to the constructed instance of the Hiding Leaders problem.

Finally, we will prove that if there exists a solution to the constructed instance of the Hiding Leaders problem, then there also exists a solution to the given instance of the 3-Set-Cover problem. We will prove this by contradiction. Assume that there exists no solution to the given instance of the 3-Set-Cover problem, but there exists a solution A^* to the constructed instance of the Hiding Leaders problem. We will show that A^* cannot be a solution, hence the contradiction. First, since there are no solutions to the 3-Set-Cover problem instance, then some nodes remain uncovered by the sets corresponding to the nodes that belong to A^* , hence $|U_{A^*}| \leq |U| - 1$. Consequently, given that $|A^*| \leq |S|$, we have: $D(G', z) - D(G', t) \leq \alpha - \beta - \gamma$. By substituting values of α , β , and γ , we get: $D(G', z) - D(G', t) \leq -1 < 0$. Therefore, we have $D(G', t) > D(G', z)$. However, we know that

$D(G', t) < D(G', z)$ is a necessary condition for A^* to be a solution to the constructed instance of the problem of Hiding Leaders. Hence, A^* cannot be such a solution.

We have shown that a solution to the constructed instance of the Hiding Leaders problem exists if and only if there also exists a solution to the given instance of the 3-Set-Cover problem, which concludes the proof. \square

Finally, we prove the NP-completeness of the problem given the betweenness centrality.

THEOREM 3. *The problem of Hiding Leaders is NP-complete given the betweenness centrality measure.*

PROOF. The problem is trivially in NP, since after the addition of a given set of edges A^* , it is possible to compute the betweenness centrality for all nodes in polynomial time.

Next, we prove that the problem is NP-hard. To this end, we propose a reduction from the NP-complete 3-Set-Cover problem. The decision version of this problem is defined by a universe $U = \{u_1, \dots, u_{|U|}\}$ and a collection of sets $S = \{S_1, \dots, S_{|S|}\}$ such that $\forall_i S_i \subset U \wedge |S_i| = 3$, where the goal is to determine whether there exist b elements of S the union of which equals U .

First, given an instance of the 3-Set-Cover problem, let us construct a network G as follows (notice that this is the same construction as in the proof of Theorem 2, but with different values of α , β , and γ):

- **The set of nodes:** For every $S_i \in S$, we create a single node denoted by S_i . For every $u_i \in U$, we create γ nodes denoted by $u'_{i,1}, \dots, u'_{i,\gamma}$, where $\gamma = |U|(|S| + 1)^2$. We denote the set of all S_i nodes by S , and the set of all $u'_{i,j}$ nodes by U' . In addition, we create α nodes denoted by $X = \{x_1, \dots, x_\alpha\}$, where $\alpha = |U|^3(|S| + 1)^2$ and create β nodes denoted by $Y = \{y_1, \dots, y_\beta\}$, where $\beta = |U|^3(|S| + 1)^2 - |U|(|S| + 1)$. Last, we create two additional nodes, denoted by t and z .
- **The set of edges:** First, we create the edge (t, z) . Then, for every node x_i , we create an edge (x_i, t) , and for every node y_i , we create an edge (y_i, z) . For every node $S_i \in S$, we create an edge (S_i, z) . For every node $u'_{j,k}$, we create an edge $(S_i, u'_{j,k})$ if and only if $u_j \in S_i$. We also create edges such that the nodes in X form a clique, i.e., we create an edge (x_i, x_j) for every $x_i, x_j \in X$. Likewise, we create edges such that the nodes in Y form a clique, i.e., we create an edge (y_i, y_j) for every $y_i, y_j \in Y$. Finally, we create edges such that the nodes in S form a clique, i.e., we create an edge (S_i, S_j) for every $S_i, S_j \in S$.

An example of the resulting network, G , is illustrated in Figure 2. Now, consider instance (G, L, b, c, δ) of the problem of Hiding Leaders, where:

- G is the network we just constructed;
- $L = \{z\} \cup U' \cup X \cup Y$;
- b is the parameter of the 3-Set-Cover problem (where the goal is to determine whether there exist b elements of S the union of which equals U);
- c is the betweenness centrality measure;
- $\delta = 1$.

From the definition of the problem of Hiding Leaders, we see that the only edges that can be added to the graph are the edges between t and the members of S , i.e., $A^* \subseteq \hat{A}$, where $\hat{A} = \{(t, S_1), \dots, (t, S_k)\}$. Notice that any such choice of A^* corresponds to selecting a subset of $|A^*|$ elements of S in the 3-Set-Cover problem. In what follows, we will show that a solution to the above instance of Hiding Leaders corresponds to a solution to the 3-Set-Cover problem.

First, we show that for every $v \in V \setminus \{t, z\}$ and every $A^* \subseteq \hat{A}$, we have $c(G', v) < c(G', t)$, where $G' = (V, E \cup A^*)$. To this end, let $B(G', v)$ denote the sum of percentages of all shortest paths between all other pairs of nodes that are *controlled* by v (following convention, a shortest path is said to be “controlled” by a node if the path goes through that node). More formally, we have: $B(G', v) = \sum_{w, w' \in V \setminus \{v\}} \frac{|\{p \in \Pi(w, w') : v \in p\}|}{|\Pi(w, w')|}$. Notice that $c(G', v) = \frac{2}{(n-1)(n-2)} B(G', v)$, hence for any two nodes $v, w \in V$, we have that $B(G', v) < B(G', w)$ implies $c(G', v) < c(G', w)$. We show that the following holds: $\forall v \in V \setminus \{t, z\} \forall A^* \subseteq \hat{A} B(G', v) < B(G', t)$.

Notice that since t controls all shortest paths between nodes in X and nodes in $\{z\} \cup Y \cup S \cup U'$ (and does not control any other shortest paths), we have that: $B(G', t) = \alpha(\beta + |U|\gamma + |S| + 1) > |U|^6(|S| + 1)^4$.

For nodes in $X \cup Y \cup U'$, we have $B(G', x_i) = B(G', y_i) = B(G', u'_{i,j}) = 0 < B(G', t)$, as they do not control any shortest paths (because all their neighbors form a clique).

For any node $S_i \in S$, we have

$$B(G', S_i) \leq 3\gamma(\alpha + \beta + (|U| - 3)\gamma + |S|) + (\alpha + 1)((|U| - 3)\gamma + |S| - 1),$$

because S_i controls some shortest paths between each of the 3γ nodes, $u'_{j,k}$, that are connected to S_i and the nodes in $\{t, z\} \cup X \cup Y \cup S \cup U' \setminus \{S_i, u'_{j,k}\}$ (there are at most $3\gamma(\alpha + \beta + (|U| - 3)\gamma + |S|)$ pairs of these nodes), and also because S_i controls some of the shortest paths between nodes in $\{t\} \cup X$ and nodes in $S \cup U' \setminus (\{S_i\} \cup (N(S_i) \cap U'))$ (there are at most $(\alpha + 1)((|U| - 3)\gamma + |S| - 1)$ pairs of these nodes). By substituting values of α , β , and γ , we get

$$B(G', S_i) \leq 3|U|(|S| + 1)^2 \left(2|U|^3(|S| + 1)^2 - |U|(|S| + 1) + (|U| - 3)|U|(|S| + 1)^2 + |S| \right) \\ + \left(|U|^3(|S| + 1)^2 + 1 \right) \left((|U| - 3)|U|(|S| + 1)^2 + |S| - 1 \right).$$

Using $|S| - |U|(|S| + 1) < 0$, $|U| - 3 < |U|$ and $|S| - 1 - 3|U|(|S| + 1)^2 < 0$, we get

$$B(G', S_i) < 3|U|(|S| + 1)^2 \left(2|U|^3(|S| + 1)^2 + |U|^2(|S| + 1)^2 \right) \\ + \left(|U|^3(|S| + 1)^2 + 1 \right) |U|^2(|S| + 1)^2.$$

Using basic arithmetic operations, this can be simplified to

$$B(G', S_i) < 3|U|^3(|S| + 1)^4(2|U| + 1) + |U|^2(|S| + 1)^2 \left(|U|^3(|S| + 1)^2 + 1 \right).$$

Since $2|U| + 1 < \frac{7}{3}|U|$ and $|U|^3(|S| + 1)^2 + 1 < 2|U|^3(|S| + 1)^2$, we get for $|U| \geq 4$:

$$B(G', S_i) < 7|U|^4(|S| + 1)^4 + 2|U|^5(|S| + 1)^4 < |U|^6(|S| + 1)^4 < B(G', t).$$

Hence, $B(G', S_i) < B(G', t)$.

Therefore, either t or z has the highest betweenness centrality. Since $z \in L$, $t \in F$, and the safety margin is $\delta = 1$, then $A^* \subseteq \hat{A}$ is a solution to the problem of Hiding Leaders if and only if $B(G', t) > B(G', z)$.

As stated above, we have that $B(G', t) = \alpha(\beta + |U|\gamma + |S| + 1)$ as t controls all shortest paths between nodes in X and all other nodes (there are $\alpha(\beta + |U|\gamma + |S| + 1)$ such pairs of nodes) and does not control any other shortest paths.

However, we have that

$$B(G', z) = \beta(\alpha + |U|\gamma + |S| + 1) + \frac{(\alpha + 1)(|S| - |A^*|)}{|A^*| + 1} + \sum_{\substack{u'_{i,j} \in U' \\ N(t, u'_{i,j})=0}} \frac{(\alpha + 1)|N(z, u'_{i,j})|}{|N(z, u'_{i,j})| + |A^*||N(z, u'_{i,j})|},$$

as z controls all shortest paths between nodes in Y and all other nodes (there are $\beta(\alpha + |U|\gamma + |S| + 1)$ such pairs of nodes), one of $|A^*| + 1$ shortest paths between each node in $\{t\} \cup X$ and each of $|S| - |A^*|$ nodes in $S \setminus \{S_i : (t, S_i) \in A^*\}$, and $|N(z, u'_{i,j})|$ of the shortest paths between $\{t\} \cup X$ and nodes $\{u'_{i,j} \in U' : N(t, u'_{i,j}) = \emptyset\}$. Therefore, we have that

$$B(G', z) - B(G', t) = (\beta - \alpha)(|U|\gamma + |S| + 1) + (\alpha + 1) \left(\frac{|S| - |A^*|}{|A^*| + 1} + \sum_{\substack{u'_{i,j} \in U': \\ N(t, u'_{i,j}) = \emptyset}} \frac{1}{|A^*| + 1} \right).$$

We will now prove that if there exists a solution S^* to the given instance of the 3-Set-Cover problem, then there also exists a solution to the constructed instance of the Hiding Leaders problem. To this end, let $A^* = \{(t, S_i) : S_i \in S^*\}$. We will show that A^* is a solution to the constructed instance of the Hiding Leaders problem. Since S^* is a solution to the given instance of the 3-Set-Cover problem, then in $G' = (V, E \cup A^*)$, we have that for every node $u'_{j,k}$ there exists a node S_i such that $S_i \in N(t, u'_{j,k})$ (it is a node $S_i \in S^*$ such that $u_j \in S_i$). Hence, given that $|A^*| > 0$, we have

$$B(G', z) - B(G', t) < (\beta - \alpha)(|U|\gamma + |S| + 1) + (\alpha + 1)|S|.$$

By substituting values of α , β , and γ , we get

$$B(G', z) - B(G', t) < (|U|^3(|S| + 1)^2 + 1)|S| - |U|(|S| + 1)(|U|^2(|S| + 1)^2 + |S| + 1).$$

Using basic arithmetic operations, we get

$$B(G', z) - B(G', t) < (|U|^3(|S| + 1)^2 + 1)|S| - (|U|^3(|S| + 1) + |U|)(|S| + 1)^2 < 0.$$

Therefore, we have $B(G', t) > B(G', z)$, which, as mentioned above, is a sufficient condition for A^* to be a solution to the constructed instance of the Hiding Leaders problem.

Finally, we will prove that if there exists a solution to the constructed instance of the Hiding Leaders problem, then there also exists a solution to the given instance of the 3-Set-Cover problem. We will prove this by contradiction. Assume that there exists no solution to the given instance of the 3-Set-Cover problem, but there exists a solution A^* to the constructed instance of the Hiding Leaders problem. We will show that A^* cannot be a solution, hence the contradiction.

Since there are no solutions to the 3-Set-Cover problem instance, after adding A^* to the network, there must exist a node $u'_{i,j}$ such that $N(t, u'_{i,j}) = \emptyset$, because otherwise the set $\{S_i \in S : (t, S_i) \in A^*\}$ would be such a solution. Since for a given i all nodes $u'_{i,1}, \dots, u'_{i,\gamma}$ have the same set of neighbors, there must exist at least γ nodes $u'_{i,j}$ such that $N(t, u'_{i,j}) = \emptyset$. Hence, given that $|A^*| \leq |S|$, we have

$$B(G', z) - B(G', t) \geq (\beta - \alpha)(|U|\gamma + |S| + 1) + (\alpha + 1)\gamma \frac{1}{|S| + 1}.$$

By substituting values of α , β , and γ , we get

$$B(G', z) - B(G', t) \geq |U|(|S| + 1)(|U|^3(|S| + 1)^2 + 1) - |U|(|S| + 1)(|U|^2(|S| + 1)^2 + |S| + 1).$$

Using basic arithmetic operations, we get

$$B(G', z) - B(G', t) \geq |U|^4(|S| + 1)^3 + 1 - (|U|^3(|S| + 1)^3 + |U|(|S| + 1)^2) > 0.$$

Therefore, we have $B(G', t) < B(G', z)$. However, we know that $B(G', t) > B(G', z)$ is a necessary condition for A^* to be a solution to the constructed instance of the Hiding Leaders problem. Hence, A^* cannot be such a solution.

We have shown that a solution to the constructed instance of the Hiding Leaders problem exists if and only if there also exists a solution to the given instance of the 3-Set-Cover problem, which concludes the proof. \square

3.2 Hardness of Approximation

Let us now consider a minimization version of the problem of Hiding Leaders.

Definition 2 (Minimum Hiding Leaders). This problem is defined by a tuple, (G, L, c, δ) , where $G = (V, E) \in \mathbb{G}$ is a network, $L \subset V$ is a set of leaders to be hidden, $c : \mathbb{G} \times V \rightarrow \mathbb{R}$ is a centrality measure, and $\delta \in \mathbb{N}$ is a chosen safety margin. Then, if we denote by $F = V \setminus L$ the set of “followers,” then the goal is to identify a set of edges to be added to the network, $A^* \subseteq F \times F$, such that the size of A^* is minimal and the resulting network $G' = (V, E \cup A^*)$ contains at least d followers that each have a centrality score higher than that of any leader, i.e.,

$$\exists_{F' \subseteq F} |F'| \geq \delta \wedge \forall_{f \in F'} \forall_{l \in L} c(G', f) > c(G', l).$$

Intuitively, the goal is to identify the smallest possible set of connections between the followers, the addition of which will keep the leaders safe from detection by the given centrality measure.

Regarding the approximation of the problem given the degree centrality, Dey and Medya [17] presented a 2-approximation algorithm, and they showed that if there exists a $(2 - \epsilon)$ -approximation algorithm for $\epsilon > 0$, then there also exists a $\frac{\epsilon}{2}$ -approximation algorithm for the Densest k-Subgraph problem, which would be considered a substantial breakthrough.

We now give the hardness of approximation results for closeness and betweenness centrality.

THEOREM 4. *The Minimum Hiding Leaders problem given the closeness centrality cannot be approximated within a ratio of $(1 - \epsilon) \ln |F|$ for any $\epsilon > 0$, unless $P = NP$.*

PROOF. To prove the theorem, we will use the result by Dinur and Steurer [19] that the *Minimum 3-Set-Cover* problem cannot be approximated within $(1 - \epsilon) \ln |F|$ for any $\epsilon > 0$, unless $P = NP$. We will show that the existence of an efficient approximation algorithm for the Minimum Hiding Leaders problem implies the existence of an approximation algorithm for the Minimum 3-Set-Cover problem with the same approximation ratio.

Let $I = (U, S)$ be an instance of the Minimum Set-Cover problem, where U is the universe $\{u_1, \dots, u_{|U|}\}$, while S is a collection $\{S_1, \dots, S_{|S|}\}$ of subsets of U such that $\forall_{S_i} |S_i| = 3$. The goal here is to find a subset $S^* \subseteq S$ such that the union of S^* equals U and the size of S^* is minimal.

First, we will show a function $f(I)$ that, based on an instance of the problem of Minimum 3-Set-Cover, constructs an instance of the Minimum Leaders Problem. Let:

- G be a network constructed as in the proof of Theorem 2,
- $L = \{z\} \cup X \cup Y \cup U'$ be the set of leaders;
- c be the closeness centrality measure;
- $\delta = 1$ be the safety margin.

The formula of the function f is then $f(I) = (G, L, c, \delta)$. Now let A^* be the solution to the instance $f(I)$ of the Minimum Hiding Leaders problem. The function g that computes the corresponding solution to the instance I of the Minimum 3-Set-Cover problem is then $g(A^*) = \{S_i \in S : (t, S_i) \in A^*\}$, i.e., the sets corresponding to the nodes that got connected with t in the solution to $f(I)$.

In the proof of Theorem 2, we showed that A^* is a solution to the constructed instance of the Hiding Leaders problem if and only if $U_{A^*} = U$, i.e., all elements of U are covered by sets in $\{S_i \in S : (t, S_i) \in A^*\}$. Since the Minimum Hiding Leaders problem has the same set of constraints (other than $|A^*| = b$) as the Hiding Leaders problem, it is also the case that A^* is a solution to the

constructed instance of the Minimum Hiding Leaders problem if and only if all elements of U are covered by sets in $g(A^*) = \{S_i \in S : (t, S_i) \in A^*\}$. What is more, we have that $|A^*| = |g(A^*)|$.

Now, assume that there exists an approximation algorithm for the Minimum Hiding Leaders problem with ratio $(1 - \epsilon) \ln |F|$ for some $\epsilon > 0$. Let us use this algorithm to solve the constructed instance $f(I)$, acquiring solution A^* , and consider the solution $g(A^*)$ to the given instance of the Minimum 3-Set-Cover problem. Since the size of the solution is the same for both instances and $|S| = |F| - 1$, we obtained an approximation algorithm that solves the Minimum 3-Set-Cover problem to within $(1 - \epsilon) \ln |S|$ for $\epsilon > 0$. However, Dinur and Steurer [19] showed that the Minimum 3-Set-Cover problem cannot be approximated to within $(1 - \epsilon) \ln |S|$ for any $\epsilon > 0$, unless $P = NP$. Therefore, such an approximation algorithm for the Minimum Hiding Leaders problem cannot exist, unless $P = NP$. This concludes the proof. \square

THEOREM 5. *The Minimum Hiding Leaders problem given the betweenness centrality cannot be approximated within a ratio of $(1 - \epsilon) \ln |F|$ for any $\epsilon > 0$, unless $P = NP$.*

PROOF. To prove the theorem, we will use the result by Dinur and Steurer [19] that the *Minimum 3-Set-Cover* problem cannot be approximated within $(1 - \epsilon) \ln |F|$ for any $\epsilon > 0$, unless $P = NP$. We will show that the existence of an efficient approximation algorithm for the Minimum Hiding Leaders problem implies the existence of an approximation algorithm for the Minimum 3-Set-Cover problem with the same approximation ratio.

Let $I = (U, S)$ be an instance of the Minimum Set-Cover problem, where U is the universe $\{u_1, \dots, u_{|U|}\}$, while S is a collection $\{S_1, \dots, S_{|S|}\}$ of subsets of U such that $\forall S_i |S_i| = 3$. The goal here is to find a subset $S^* \subseteq S$ such that the union of S^* equals U and the size of S^* is minimal.

First, we will show a function $f(I)$ that, based on an instance of the Minimum 3-Set-Cover problem, constructs an instance of the Minimum Leaders Problem. Let:

- G be a network constructed as in the proof of Theorem 3,
- $L = \{z\} \cup X \cup Y \cup U'$ be the set of leaders;
- c be the betweenness centrality measure;
- $\delta = 1$ be the safety margin.

The formula of the function f is then $f(I) = (G, L, c, \delta)$. Now let A^* be a solution to the instance $f(I)$ of the Minimum Hiding Leaders problem. The function g that computes the corresponding solution to the instance I of the Minimum 3-Set-Cover problem is then $g(A^*) = \{S_i \in S : (t, S_i) \in A^*\}$, i.e., the sets corresponding to the nodes that got connected with t in the solution to $f(I)$.

In the proof of Theorem 3, we showed that A^* is a solution to the constructed instance of the Hiding Leaders problem if and only if there does not exist a node $u'_{i,j}$ such that $N(t, u'_{i,j}) = \emptyset$. This is the case when all elements of U are covered by sets in $\{S_i \in S : (t, S_i) \in A^*\}$. Since the Minimum Hiding Leaders problem has the same set of constraints (other than $|A^*| = b$) as the Hiding Leaders problem, it is also the case that A^* is a solution to the constructed instance of the Minimum Hiding Leaders problem if and only if all the elements of U are covered by sets in $g(A^*) = \{S_i \in S : (t, S_i) \in A^*\}$. What is more, we have that $|A^*| = |g(A^*)|$.

Now, assume that there exists an approximation algorithm for the Minimum Hiding Leaders problem with ratio $(1 - \epsilon) \ln |F|$ for some $\epsilon > 0$. Let us use this algorithm to solve the constructed instance $f(I)$, acquiring solution A^* , and consider the solution $g(A^*)$ to the given instance of the Minimum 3-Set-Cover problem. Since the size of the solution is the same for both instances and $|S| = |F| - 1$, we obtained an approximation algorithm that solves the Minimum 3-Set-Cover problem to within $(1 - \epsilon) \ln |S|$ for $\epsilon > 0$. However, Dinur and Steurer [19] showed that the Minimum 3-Set-Cover problem cannot be approximated to within $(1 - \epsilon) \ln |S|$ for any $\epsilon > 0$, unless $P = NP$.

Therefore, such an approximation algorithm for the Minimum Hiding Leaders problem cannot exist, unless $P = NP$. This concludes the proof. \square

To summarize, in this section, we have analyzed the problem of hiding leaders in a given network. This problem involves determining an optimal set of edges to be added between the followers to reduce the ranking of the leaders according to the three main centrality measures, namely, degree, closeness, and betweenness. Our theoretical findings can be summarized as follows. In Theorems 1–3, we have shown that an optimal solution cannot be computed in polynomial time given any of the centrality measures. In other words, if the followers wanted to create fake connections between them to mislead the analysis and keep the identity of their leaders hidden from centrality measures, then it is intractable to identify an optimal set of such connections. Moreover, in Theorems 4 and 5, we have shown that the problem cannot even be approximated to a sub-logarithmic ratio (with respect to the number of followers) given closeness and betweenness centrality. In other words, if the followers wanted to create fake connections to shield their leaders from closeness and betweenness centrality, then it is intractable to identify a set of fake connections whose size is relatively close to optimal (as mentioned earlier, the hardness of approximation for the degree centrality was shown by Dey and Medya [17]).

In the following section, we shift our attention to a setting where, instead of modifying an existing network, the goal is to build a network structure from scratch, such that the leaders are guaranteed to be hidden at least to a certain extent.

4 CAPTAIN NETWORK

In the previous section, we proved the NP-completeness of modifying an existing network to hide its leaders. However, in certain cases, the leaders are to develop a new covert network (e.g. a subnetwork in a foreign country) rather than to modify an existing one. In this section, we show that it is possible to efficiently create a network from scratch, designed specifically to hide its leaders without limiting their ability to influence the other nodes in the network.

4.1 Isolated Structure

We begin with a network analyzed in isolation, without being embedded in a larger structure. We call this the “captain” network, and it is formed in the following manner. First, the leader nodes, L , form a clique, to provide the best possible communication among them. Each leader $l_i \in L$ is then assigned a group of k “captains,” $C_i = \{c_{i,1}, \dots, c_{i,k}\}$, which are connected to that leader. All captains are then connected into a complete $|L|$ -partite graph based on the division into groups. We will denote the set of all captains by C , i.e., $C = \bigcup C_i$. A captain, $c_{i,j}$ serves two purposes: the first is to conceal the leaders in L , which is achieved by ensuring that $c_{i,j}$ is ranked higher than every node in L (according to the three standard centrality measures); the second purpose of $c_{i,j}$ is to pass on the influence of l_i to the rest of the network. The remaining m nodes, $X = \{x_1, \dots, x_m\}$, are each connected to one captain from each group. Note that the set of “followers” in this network is $F = C \cup X$. Algorithm 1 summarizes the steps that create such a captain network, whereas Figure 3 illustrates a sample network with $|L| = 3$.

Note that if the steps of Algorithm 1 are followed given just a single leader, the result would be a tree structure. While a tree is a fairly common organizational structure, it may not provide an adequate disguise of the leader, especially if the leader is identified as a root of the tree. With this in mind, whenever there is a single leader, we create two groups of captains to avoid the tree structure. The resulting structure is illustrated in Figure 4.

Figure 5 presents the time necessary to generate the structure of a captain network using Algorithm 1. The computation was performed on a PC with Intel Core i7-4790 CPU and 24 GB RAM,

ALGORITHM 1: Constructing a captain network with multiple leaders

Input: The set of leaders $L = \{l_1, \dots, l_{|L|}\}$, the set of followers $F = \{f_1, \dots, f_{|F|}\}$, and the number of captains in each group, i.e., k (where $1 \leq k \leq \lfloor \frac{|F|}{|L|} \rfloor$).

Output: The captain network $(L \cup F, E)$

```

for  $i = 1, \dots, \max(2, |L|)$  do
  for  $j = 1, \dots, k$  do
     $c_{i,j} \leftarrow f_{(i-1)k+j}$ 
     $C_i \leftarrow C_i \cup \{c_{i,j}\}$ 
 $X \leftarrow F \setminus \bigcup_i C_i$ 
for  $l_i, l_j \in L$  do
   $E \leftarrow E \cup \{(l_i, l_j)\}$ 
for  $l_i \in L$  do
  for  $c_{i,j} \in C_i$  do
     $E \leftarrow E \cup \{(l_i, c_{i,j})\}$ 
if  $|L| = 1$  then
  for  $c_{2,j} \in C_2$  do
     $E \leftarrow E \cup \{(l_1, c_{2,j})\}$ 
for  $C_i \neq C_j$  do
  for  $c \in C_i$  do
    for  $c' \in C_j$  do
       $E \leftarrow E \cup \{(c, c')\}$ 
 $j \leftarrow 0$ 
for  $x \in X$  do
  for  $i = 1, \dots, \max(2, |L|)$  do
     $E \leftarrow E \cup \{(x, c_{i,j})\}$ 
     $j \leftarrow (j + 1) \bmod k$ 
return  $(L \cup F, E)$ 

```

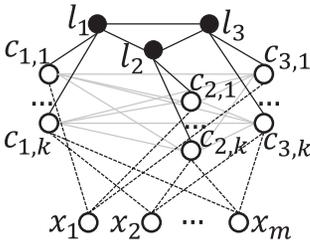


Fig. 3. A captain network with $|L| = 3$. Edges that involve the leaders are depicted as solid black lines; edges between captains are depicted as gray lines; edges between captains and other nodes are depicted as dotted lines.

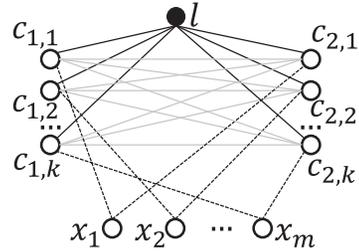


Fig. 4. A captain network with a single leader. Edges representing the leader's links are depicted as solid black lines; edges between captains are depicted as gray lines; edges between captains and other nodes are depicted as dotted lines.

using an implementation in Java programming language, JDK version 13.0.2. As it can be seen, the captain network structure can be generated efficiently even if it incorporates a million nodes.

Next, we prove that every captain has a greater centrality value than any of the leaders.

THEOREM 6. *Given a captain network with k captains in each group, let $r = \lfloor \frac{|X|}{k} \rfloor$ denote the minimal number of connections that a captain, $c_{i,j}$, has with nodes from X . Then, if either*

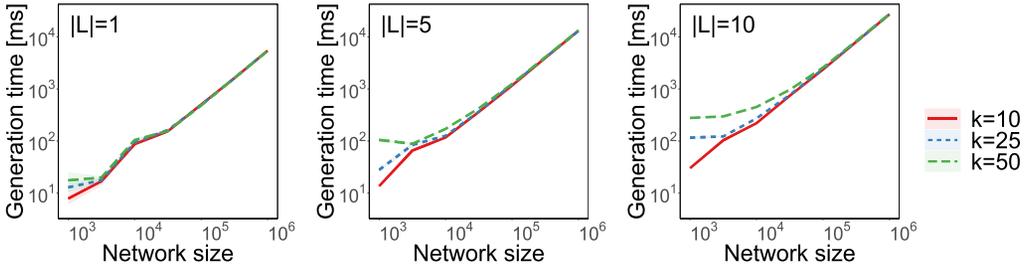


Fig. 5. Given the number of leaders, the number of captains, and the size of the captain network, the figure depicts the time it takes to generate network structure. The results are reported on logarithmic scales, as an average over 100 executions, with colored areas representing 95% confidence intervals.

($|L| \geq 2$ and $r \geq 1$), or ($|L| = 1$ and $k < \sqrt{|F| + 1} - 1$), then every captain has a greater degree, closeness, and betweenness centrality than any of the leaders.

PROOF. Starting with the case of degree centrality and multiple leaders, the degree of a leader node, l , is $c_{dg}(G, l) = \frac{|L|+k-1}{n-1}$, since it is only connected to other leaders and captains from its group. However, the degree of a captain, $c_{i,j}$, is $c_{dg}(G, c_{i,j}) \geq \frac{1+k(|L|-1)+r}{n-1}$, since it is connected to one of the leader nodes, to all captains from other groups, and to at least r other nodes from X . As such, we have

$$c_{dg}(G, c_{i,j}) - c_{dg}(G, l) \geq \frac{1 + k(|L| - 1) + r - (|L| + k - 1)}{n - 1} = \frac{(|L| - 2)(k - 1) + r}{n - 1}.$$

Now, since $|L| \geq 2$, $k \geq 1$, and $r \geq 1$, then we have that $c_{dg}(G, c_{i,j}) > c_{dg}(G, l)$ for any $c_{i,j}$.

Next, we handle the case of degree centrality with a single leader. The degree of the leader node, l , is $c_{dg}(G, l) = \frac{2k}{n-1}$, since it is only connected to captains from both groups. However, the degree of a captain $c_{i,j}$, is $c_{dg}(G, c_{i,j}) \geq \frac{1+k+r}{n-1}$, since it is connected to the leader node, to all captains from other groups, and to at least r members. Thus: $c_{dg}(G, c_{i,j}) - c_{dg}(G, l) \geq \frac{1+k+r-2k}{n-1} = \frac{r-k+1}{n-1}$. Therefore, since $r = \lfloor \frac{|X|}{k} \rfloor$, $|X| = |F| - 2k$, then we have $c_{dg}(G, c_{i,j}) > c_{dg}(G, l)$ for: $\lfloor \frac{|F|-2k}{k} \rfloor - k + 1 \geq \frac{|F|-2k}{k} - k > 0$, which is the case for $k < \sqrt{|F| + 1} - 1$.

Moving on to closeness centrality, for any given node, v , this centrality depends inversely on the sum of the lengths of the shortest paths from v to every other node, i.e., $\sum_{w \in V} d(v, w)$. For every leader and every captain, the distance to every other node is either 1 or 2. More precisely, for every $v \in L \cup C$, we have: $\sum_{w \in V} d(v, w) = 1|N(v)| + 2(n - |N(v)|) = 2n - |N(v)|$. Consequently, whenever all captains have greater degree centrality than all leaders, they must also have greater closeness centrality. Since we have already proven this fact for the degree centrality, then this implies that $c_{cl}(G, c_{i,j}) > c_{cl}(G, l)$.

Finally, regarding betweenness centrality, let $\zeta(v)$ denote: $\sum_{u, w \in V \setminus \{v\}} \frac{|\{p \in \Pi(u, w) : v \in p\}|}{|\Pi(u, w)|}$. Then the betweenness centrality of a node $v \in V$ can be written as $c_{bt}(G, v) = \frac{2}{(n-1)(n-2)} \zeta(v)$.

For a network with multiple leaders, every leader node l belongs to one of the $(|L| - 1)k + 1$ shortest paths between pairs of captains from her group (alternative shortest paths run through captains from other groups), as well as one of the $k + 1$ shortest paths between each captain from her group and all other leaders (alternative shortest paths run through captains from the group of the chosen leader). Since the leader node l belongs to no other shortest paths, we have: $\zeta(l) = \frac{k(k-1)}{2((|L|-1)k+1)} + \frac{k(|L|-1)}{k+1}$. Having analyzed $\zeta(l)$, let us now analyze $\zeta(c_{i,j})$ for a captain, $c_{i,j}$. In particular, since $c_{i,j}$ belongs to one of the $(|L| - 1)k + 1$ shortest paths between pairs of captains

from all other groups, as well as one of the $k + 1$ shortest paths between each captain from other groups and the leader of her group, we have

$$\zeta(c_{i,j}) > \frac{(|L| - 1)k(k - 1)}{2((|L| - 1)k + 1)} + \frac{k(|L| - 1)}{k + 1} > \frac{k(k - 1)}{2((|L| - 1)k + 1)} + \frac{k(|L| - 1)}{k + 1} = \zeta(l).$$

Notice that there also exist other shortest paths controlled by a given captain $c_{i,j}$ (e.g., shortest paths between any two nodes $x_{i'}$ and $x_{j'}$ connected to $c_{i,j}$). However, they are not required to show the inequality; hence, we do not consider them. We have that $\zeta(c_{i,j}) > \zeta(l)$, which implies that $c_{bt}(G, c_{i,j}) > c_{bt}(G, l)$.

For a network with a single leader, the leader node l belongs to one of the $k + 1$ shortest paths between pairs of captains from each group (alternative shortest paths run through captains from other groups). Since the leader node l belongs to no other shortest paths, we have: $\zeta(l) = \frac{k(k-1)}{k+1}$. Having analyzed $\zeta(l)$, let us now analyze $\zeta(c_{i,j})$ for a captain, $c_{i,j}$. In particular, since $c_{i,j}$ belongs to one of the $k + 1$ shortest paths between pairs of captains from the other group, and to the only shortest path between member nodes connected to her and captains from the other group, we have

$$\zeta(c_{i,j}) \geq \frac{k(k - 1)}{2(k + 1)} + r(k - 1) = \frac{k(k - 1) + 2r(k - 1)(k + 1)}{2(k + 1)} > \frac{k(k - 1)}{k + 1} = \zeta(l).$$

Therefore, we have that $\zeta(c_{i,j}) > \zeta(l)$, which implies that $c_{bt}(G, c_{i,j}) > c_{bt}(G, l)$. \square

As for the follower's centrality values, they are sure to be lower than those of the captain. Intuitively, the follower has fewer neighbors than the captain to which she is connected, and she lies on none of the shortest paths in the network, as her neighbors form a clique.

4.2 Captain Network Embedded in a Larger Structure

Arguably, it might be too strict to have a network structure in which ordinary members of the organization, i.e., the nodes in X , cannot communicate directly with each other. Furthermore, the captain network might not be analyzed in isolation, but rather as part of a larger social network in which it is embedded. The following analysis takes these issues into consideration. Specifically, we start by considering a situation in which the members of X can have any structure of connections between them, and they can even contact people from outside the organization. After that, we consider a similar situation, but where the captain is also allowed to contact people from outside the organization.

THEOREM 7. *Let G be a captain network such that $|L| \geq 2$ and $\lfloor \frac{|X|}{k} \rfloor \geq 1$ and let $G' = (V', E')$ be another (potentially empty) network. After adding any set of edges $A^* \subseteq X \times (X \cup V')$, every captain has greater degree, closeness and betweenness centrality than any of the leaders.*

PROOF. As for degree centrality, adding A^* does not change the degree of any of the leaders nor the captains in the network. Thus, the captains still have a higher degree than the leaders.

As for closeness centrality, notice that the distance between any leader and any node in X is always 2. Moreover, the distance between any captain and any node in X is not greater than 2. Hence, we have that the distance from any leader to a given node in X is greater than or equal to the distance from any captain to the same node in X , i.e., $\forall l \in L \forall c_{i,j} \in C \forall x \in X d(l, x) \geq d(c_{i,j}, x)$. Since all paths from the nodes in $L \cup C$ to the nodes in V' run through the nodes in X , then the same holds for the nodes in V' , i.e., $\forall l \in L \forall c_{i,j} \in C \forall v \in V' d(l, v) \geq d(c_{i,j}, v)$. As for the shortest paths between the nodes in $L \cup C$, none of them run through the nodes in $X \cup V'$, neither before, nor after the addition of A^* . Therefore, since for the network G we know that every captain has greater closeness centrality than any leader, this would still hold after the addition of A^* .

Finally, regarding betweenness centrality, notice that when we compared the leaders to the captains in a network with multiple leaders in the proof of Theorem 6, we only considered the paths between the nodes in $L \cup C$. Since, before the addition of A^* , none of the shortest paths between the nodes in $L \cup C$ go through the nodes in X , and since the addition of A^* does not create any new shortest paths between the nodes in $L \cup C$, then the argument made in the proof of Theorem 6 about $c_{bt}(G, c_{i,j}) > c_{bt}(G, l)$ remains valid. Notice that this does not mean that the betweenness centrality of captains does not change; it only means that their betweenness centrality remains greater than that of the leaders. \square

It is worth noticing that for a captain network with only one leader the result holds only for degree and closeness centrality measures (given the assumptions of Theorem 6), i.e., adding new connections can cause the leader to have higher betweenness centrality than the captains. For example, connecting the members of X into a clique would create alternative shortest paths (running through other nodes in X) between any node x_i and the captains that x_i is not directly connected to, thus greatly reducing the betweenness centrality of the captains that x_i is directly connected to (and who previously controlled all shortest paths between x_i and captains from the other group). At the same time, the betweenness centrality of the sole leader remains unaffected by adding connections between the members of X . One possible solution in such a situation is to choose one of the ordinary members as an *ad hoc* leader and build the structure for $|L| = 2$.

THEOREM 8. *Let $G = (V, E)$ be a captain network such that $|L| \geq 2$ and $\lfloor \frac{|X|}{k} \rfloor \geq 1$ and let $H = (V', E)$ be another (potentially empty) network. After adding any set of edges $A^* \subseteq (C \times V') \cup (X \times (X \cup V'))$:*

- (1) every captain has greater degree centrality than any of the leader nodes;
- (2) for every leader, at least $k(|L| - 1)$ captains have greater closeness centrality than the leader.

PROOF. As for degree centrality, adding A^* does not change the degree of the leaders, and it can only increase the degree of the captains. Consequently, the captains still have a higher degree than the leaders.

As for closeness centrality, we will show that for any given leader l_i , every captain from a group other than her own, i.e., any $c_{j,k}$ such that $j \neq i$, has greater closeness centrality than l_i . To this end, we show that the following holds, where $D(G^*, v) = \frac{n-1}{c_{cl}(G^*, v)} = \sum_{w \in V \cup V'} d_{G^*}(v, w)$ and G^* is the union of G and G' after the addition of A^* : $\forall l_i \in L \forall c_{j,k} \in C: j \neq i D(G^*, l_i) > D(G^*, c_{j,k})$. We have that: $D(G^*, l_i) = |L| - 1 + k + 2(|L| - 1)k + 2|X| + \sum_{v' \in V'} d(l_i, v')$ as l_i is at distance 1 from the other $|L| - 1$ leaders, as well as k captains from her group, and at distance 2 from the other $(|L| - 1)k$ captains and all $|X|$ members of X . We also have that

$$D(G^*, c_{j,k}) \leq 1 + (|L| - 1)k + 2(|L| - 1) + 2(k - 1) + r + 2(|X| - r) + \sum_{v' \in V'} d(c_{j,k}, v'),$$

as $c_{j,k}$ is at distance 1 from the leader of her group, all $(|L| - 1)k$ captains from other groups, as well as at least r nodes from X , and she is at distance 2 from the other $|L| - 1$ leaders, $k - 1$ captains from her own group, and to the other $|X| - r$ members of X . Hence, we have that

$$D(G^*, l_i) - D(G^*, c_{j,k}) \geq (|L| - 2)(k - 1) + r + \sum_{v' \in V'} (d(l_i, v') - d(c_{j,k}, v')).$$

We have that $|L| \geq 2$, $k \geq 1$, and $r > 0$. We also have that for a given leader l_i and for any node v' any shortest path between them must run through a node from C . If this node from C belongs to the group of l_i (in which case l_i can reach it in one step), then it can also be reached by any $c_{j,k}$ such that $j \neq i$ in one step. If this node from C belongs to another group of captains (in which case

l_i can reach it in two steps), then it can also be reached by any $c_{j,k} : j \neq i$ in at most two steps. Therefore, we have that $d(l_i, v') \geq d(c_{j,k}, v')$, which in turn implies that $D(G^*, l_i) > D(G^*, c_{j,k})$. This concludes the proof. \square

Importantly, the result does not hold for betweenness centrality. It is relatively easy to create a network where only one captain has greater betweenness centrality than a given leader. This can be achieved by connecting one of the captains from the leader's group to a single node from a large enough network H (without adding any other connections). This observation suggests that whenever there is a risk of the network being analyzed using betweenness centrality, the contacts of the captains with the outside world should be limited.

It is worth noting that the results presented in this section can be applied by leaders who do not want to build an entirely new organization, but rather modify an already existing network. In that case, instead of completely disassembling the current structure, the leaders could simply rewire the connections of themselves, and a small group of nodes selected as captains, while keeping all ties between the remaining members as they were. Such a process would still require some compromises for the sake of leaders' safety, e.g., they would no longer be able to communicate with anyone besides other leaders and some of the captains, but from the perspective of an ordinary member of the organization, such a transition could be virtually unnoticeable.

5 SIMULATION RESULTS

We now present two studies utilizing simulations to analyze different aspects of captain networks. In the first study, we investigate how modifying the parameters of the captain network affects the centrality and influence of the leaders. In the second study, we compare the resilience of the captain network to centrality-based attacks with different types of network structures.

5.1 Centrality and Influence of Leaders

As stated in Theorem 6, a captain network can indeed shield its leaders from centrality analysis based on degree, closeness and betweenness centralities. However, as far as other centrality measures and the influence of the leaders is concerned, we evaluate the network empirically. To this end, given a captain network with 100 and 1,000 nodes, we vary the parameters of the network, namely, k (the number of captains in each group) and $|L|$ (the number of leaders). For every pair of parameters, we measure the difference in terms of centrality between a leader and any given captain (the greater the difference, the greater the leaders' disguise), and we also measure the influence of a leader to see how this influence is affected by changing the parameters of the constructed network. When measuring the influence, we use either the Independent Cascade model with probability 0.15 on each edge, the Linear Threshold model with the threshold values sampled uniformly at random, or the Shapley centrality, which in the literature is often used as the measure of influence [45, 59].

Let us briefly comment on the use of Independent Cascade and Linear Threshold diffusion models as proxies for the influence of the leaders. In many cases, it is enough for the network to remain connected for the messages of the leader to reach all members (indeed, we also analyze the resilience of the captain network to be disconnected in Section 5.2). Nevertheless, in the case of covert networks, the communication channels may be monitored, interrupted, or distorted, making nondeterministic communication a more realistic model [60]. What is more, an effective communication structure can improve the cohesion of the organization, while limited access to the ideas spread by the leadership can be detrimental to the morale of the ordinary members [30]. As a result, the communication efficiency and the ability to spread ideas throughout the network are important aspects of the network structure, and diffusion models can be used to measure these

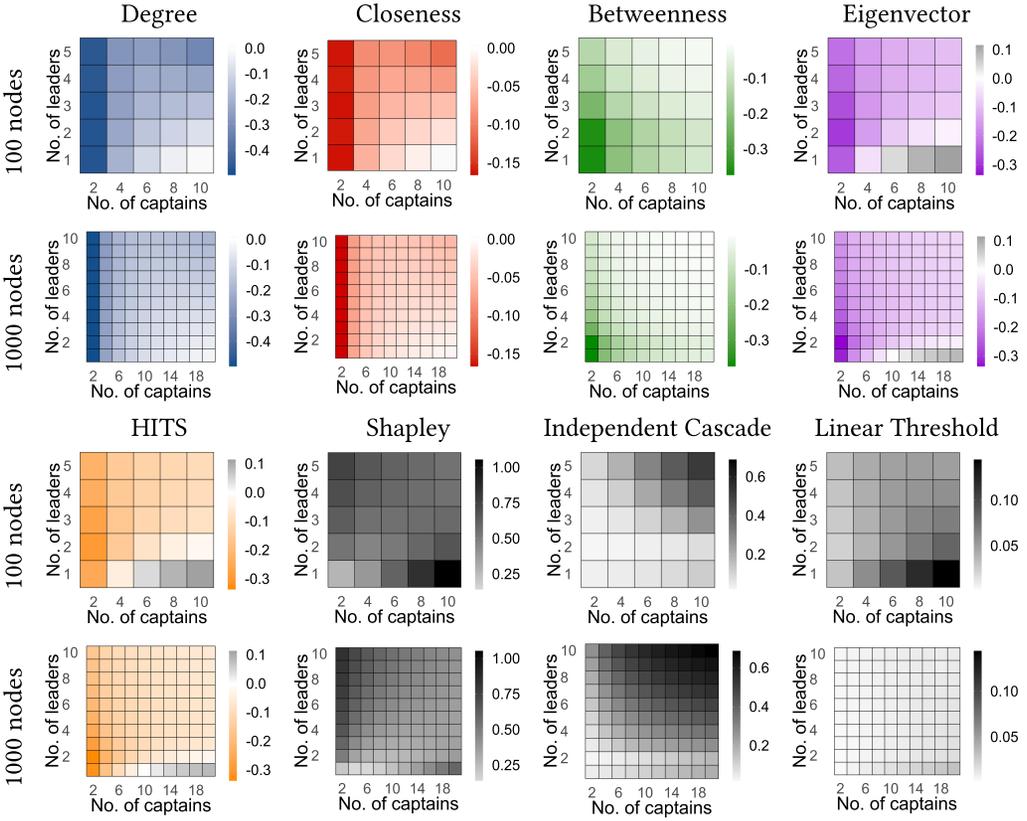


Fig. 6. Given captain networks of 100 nodes and 1,000 nodes, with varying numbers of captains per group (the x -axis) and varying numbers of leaders (the y -axis), the figure depicts the difference in terms of centrality between a leader and a captain, as well as the influence of a leader as either the Shapley centrality value or as an expected percentage of activated nodes, according to either the Independent Cascade or the Linear Threshold model. The scale is fixed for each centrality to allow for comparisons.

aspects [32, 33]. Finally, even though both models could be used as centrality measures to evaluate the importance of a given node, they do not fit this role particularly well, as they have very high computational costs compared to the centrality measures considered in this work. This is particularly problematic in a situation where we need to compute the centrality score for every single node of the network to obtain a ranking. By using them to model influence we avoid this issue, as the influence needs to be computed only for the leaders of the network.

The results are depicted in Figure 6. The x -axis represents the number of captains in each group, and the y -axis represents the number of leaders of the network. The greater the intensity of the color, the greater the difference in terms of centrality between a leader and a captain (i.e., the safer the leader), or the greater the influence of each leader.

Generally speaking, in networks with a small number of captains per group, the difference between leaders and captains in terms of centrality is typically greater than the difference when the number of captains per group is large. For networks with many captains, the difference in betweenness centrality is greater when the network has few leaders, whereas the difference in all other centrality measures is greater in the case of many leaders. As for the influence, Figure 6 shows

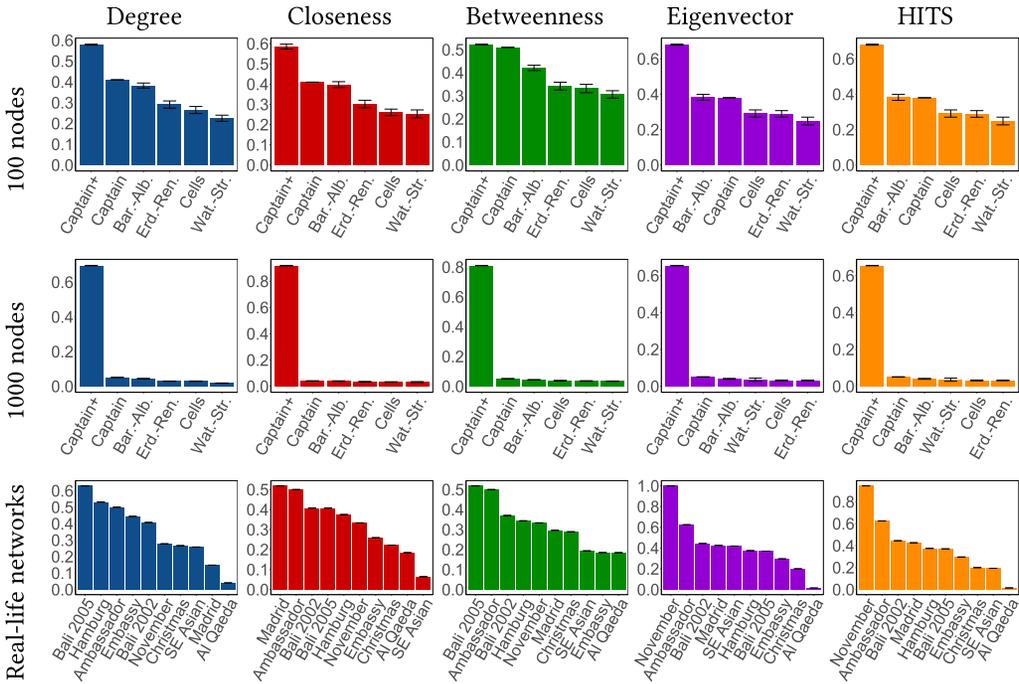


Fig. 7. Given networks of 100 nodes and 1,000 nodes, as well as real-life terrorist networks, the figure depicts the proportion of nodes that have to be removed from the network until the first leader gets removed. For random networks the results are reported as an average over 100 different networks, with error bars representing 95% confidence intervals.

that leaders can spread their messages most efficiently in a network with many captains and many leaders according to the Independent Cascade model, while for the Linear Threshold model and the Shapley centrality the best structure appears to be a network with many captains but only a few leaders.

5.2 Attack Tolerance of Captain Networks

We now present an experiment designed to evaluate the *attack tolerance* [1] of captain networks and to compare that to the attack tolerance of other network structures. To this end, we assume that a given network G is attacked by a third party utilizing a given centrality measure c . The attack is carried out by iteratively removing a node with the highest value of centrality c . We measure (i) the number of nodes that have to be removed until a leader is removed from the network, and (ii) the number of nodes that have to be removed for the network to become disconnected. If multiple nodes have the same centrality value and one of them happens to be a leader, then we assume that it is the leader that will be removed. Otherwise, we break ties uniformly at random.

We compare the captain networks containing 5 leaders and 10 captains in each group with the following random network generation models:

- Barabási-Albert networks [3], which are generated using the preferential attachment model;
- Erdős-Rényi networks [21], in which an edge is created between each pair of nodes with a constant probability;
- Watts-Strogatz networks [69] meant to represent a small-world structure, characterized with a short average distance between any pair of nodes;

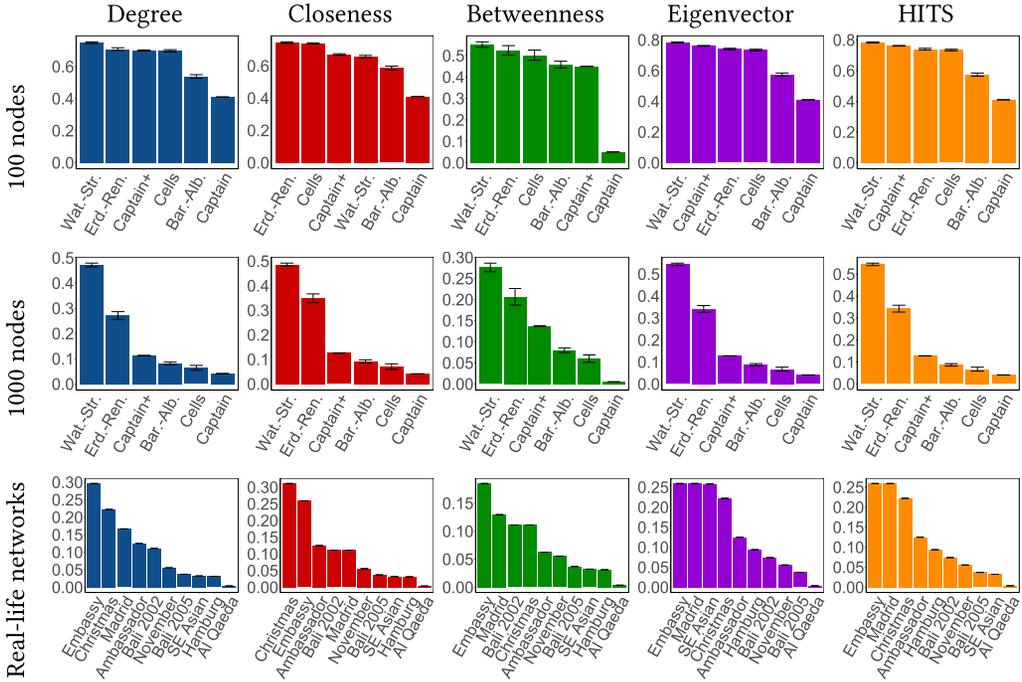


Fig. 8. Same as Figure 7, but the plots depict the proportion of nodes that have to be removed from the network until the network gets disconnected.

- Cellular networks [61] designed to recreate the structure of real-life covert networks, that often consist of several loosely connected, but very dense cells.

To keep different types of network structures comparable, we consider random networks with the same average degree as the captain network. We set the rewiring probability in the Watts-Strogatz model to $\frac{1}{4}$. The cellular network parameters are set according to Tsvetovat and Carley [61]. Additionally, we also consider a captain network where member nodes are connected into a preferential attachment network with an average degree of 50, we denote these networks as *Captain+*.

For the purpose of comparison, we also provide results of the same experiment for the following real-life terrorist networks from the ARTIS database [70]:

- Al Qaeda worldwide operations attack series 1993–2003—271 nodes, 767 edges,
- Australian embassy bombing 2004, Indonesia—27 nodes, 112 edges,
- Bali bombings 2002, Indonesia—27 nodes, 158 edges,
- Bali bombings 2005, Indonesia—27 nodes, 102 edges,
- Christmas Eve bombings 2000, Indonesia—45 nodes, 234 edges,
- November 17 organization aggregate attack series, Greece—18 nodes, 46 edges,
- the Hamburg 9/11 cell 2001, Germany—32 nodes, 121 edges,
- Madrid train bombings 2004, Spain—54 nodes, 226 edges,
- Philippines ambassador residence bombing 2000, Jakarta—16 nodes, 69 edges,
- Southeast Asian aggregate attack series 2005, Indonesia—31 nodes, 38 edges.

As most of the real-life datasets are small, we assume that they have only one leader (unlike in captain and randomly generated networks, where we assume the existence of five leaders).

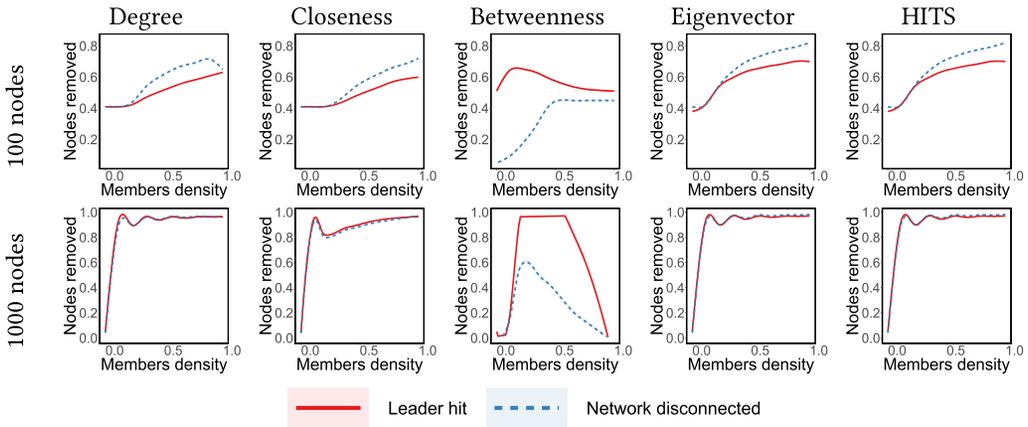


Fig. 9. Given the density of connections between the member nodes, the figure depicts the proportion of nodes that have to be removed from the network until a given effect is achieved. The results are reported as an average over 100 different networks, with colored areas representing 95% confidence intervals.

In a captain network, the leaders' identities are determined by the algorithm of network construction. For other types of networks, however, we need to decide which nodes will be considered the leaders. To be able to compare the results for random networks with the results for captain networks, we assume that the leaders initially occupy the same positions in centrality rankings. Recall that in a captain network all the captains occupy the first $k|L|$ positions in the ranking, followed by all the leaders. Based on this, for any given random network G and centrality measure c , we first rank all the nodes in G according to c , and then specify the leaders to be the nodes in positions $k|L| + 1$ to $k|L| + |L|$ in that ranking.

Figures 7 and 8 presents the results of our simulations. Each column corresponds to a scenario in which the “attacker” utilizes a specific centrality measure, while each row presents results for different network size. Each bar represents the average proportion of nodes that had to be removed from the network to either remove a leader or disconnect the network.

Real-life terrorist networks exhibit relatively high levels of resilience, although the reason may be their small size (notice how large random networks are significantly more fragile than their smaller counterparts). As can be seen, the captain network is one of the most resilient out of all considered networks of the same size in terms of protecting the leaders from being identified, regardless of the centrality measure being used. However, it is actually the easiest network structure to disconnect. This is because in a version of the captain network that contains no connections between the member nodes removing just a few captains can completely separate multiple member nodes from the network. Nevertheless, as we have shown in Section 4.2, it is possible to add any number of connections between the member nodes without sacrificing the safety of the leaders. Indeed, after adding a scale-free structure between member nodes of the captain network (bars labeled in Figures 7 and 8 as *Captain+*) the networks becomes significantly more resilient to being disconnected.

To study how the density of connections between the member nodes can affect the resilience of the network structure, we perform a series of simulations where we randomly add edges between the member nodes, while recording the number of nodes that have to be removed to hit the leader or disconnect the network. The results of our simulations are presented in Figure 9.

As can be seen, adding connections between the member nodes can significantly increase the resilience of the network both in terms of preventing the removal of leaders and in terms of keeping

the network connected. With more connections between them, member nodes are not only harder to separate from the network but also act as another buffer (in addition to the captain nodes) that protects the leaders from getting targeted. Interestingly, as shown in Figure 9, the greater density of the connections between the members does not always result in greater network resilience, especially when betweenness centrality is being used by the attacker. In very dense network structures, nodes typically have low betweenness centrality, e.g., all nodes in a clique have a betweenness centrality of zero, since the shortest path between any two nodes consists of just a single edge.

6 DISCUSSION AND CONCLUDING REMARKS

The model studied in this article offers new insights into the secrecy-efficiency tradeoff faced by covert organizations. The novelty of our approach comes from our definition of secrecy, which assumes that the members of a terrorist network act strategically to evade detection by centrality measures. Indeed, it is well established that centrality measures belong to the key social network analysis tools that are used to analyze covert networks. Unfortunately, centrality measures—like most other social network analysis tools—were designed to analyze social networks among members of the general public, rather than among adroit members of covert organizations who are well aware of the possibility of attracting unwanted attention from the authorities. However, recent findings—for, instance, with respect to ISIS—strongly suggest that such an assumption is too far-fetched [36, 50].

Our work constitutes a step towards relaxing this assumption, and contributes to the literature on the strategic analysis of social networks [47]. In particular, we showed that identifying an optimal set of edges to add to the network to decrease the leaders' ranking (according to degree, closeness, and betweenness centrality) is NP-complete. While this is a “negative” result from the computational point of view, it is in fact rather positive news for law-enforcement agencies.

The above hardness results are general in the sense that they were obtained without any considerations of the “efficiency” part of the aforementioned secrecy-efficiency tradeoff. We introduced such efficiency into the model by investigating how the leaders could construct a network from scratch so that they are adequately hidden from the three fundamental centrality measures, and adequately influential at the same time.

The network that we construct from scratch has a group of leaders forming a clique (which ensures efficient communication among them) and has a well-defined core of “captains” who are densely connected among themselves and who act as intermediaries between leaders and other members of the organization. It is known that such “inner circles” exist in some real-life terrorist networks such as, e.g., Al-Qaeda [4] and IRA [56].

Our model can be extended in various directions. First, we assume that the “evaders” (i.e., the members of the covert organization) are strategic, whereas the “seeker” (who is using centrality measures to identify key terrorist) is not, i.e., he or she is unaware of any potential strategic efforts by the evaders. It would be interesting to see new social network analysis tools, and centrality measures, in particular, that are immune (at least to some extent) against such evasion techniques.

Second, although our captain networks appear to be effective in terms of influence (i.e., they are empirically shown to grant the leaders a reasonable level of influence), they do not provide any worst-case guarantees on solution quality in this regard. This problem constitutes another direction for future research.

Finally, it would be interesting to investigate whether there exist special classes of networks for which the problem of hiding leaders can easily be solved or whether it is possible to construct a network that conceals certain edges [67].

REFERENCES

- [1] Réka Albert, Hawoong Jeong, and Albert-László Barabási. 2000. Error and attack tolerance of complex networks. *Nature* 406, 6794 (2000), 378.
- [2] Jac M. Anthonisse. 1971. The rush in a directed graph. *Stichting Mathematisch Centrum. Mathematische Besliskunde* 71, BN 9 (1971), 1–10.
- [3] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.
- [4] Victoria Barber. 2015. The evolution of Al Qaeda’s global network and Al Qaeda core’s position within it: A network analysis. *Perspect. Terror.* 9, 6 (2015), 2–35.
- [5] Alex Bavelas. 1948. A mathematical model for group structures. *Hum. Organiz.* 7, 3 (1948), 16–30.
- [6] Murray A. Beauchamp. 1965. An improved index of centrality. *Behav. Sci.* 10, 2 (1965), 161–163.
- [7] Battista Biggio, Giorgio Fumera, and Fabio Roli. 2010. Multiple classifier systems for robust classifier design in adversarial environments. *Int. J. Mach. Learn. Cybernet.* 1, 1–4 (2010), 27–41.
- [8] Aleksandar Bojchevski and Stephan Günnemann. 2019. Adversarial attacks on node embeddings via graph poisoning. In *Proceedings of the International Conference on Machine Learning (PMLR’19)*. 695–704.
- [9] Phillip Bonacich. 1987. Power and centrality: A family of measures. *Amer. J. Sociol.* 92, 5 (1987), 1170–1182.
- [10] Michael Brückner, Christian Kanzow, and Tobias Scheffer. 2012. Static prediction games for adversarial learning problems. *J. Mach. Learn. Res.* 13, (Sep. 2012), 2617–2654.
- [11] Kathleen M. Carley. 2006. Destabilization of covert networks. *Comput. Math. Organiz. Theory* 12, 1 (2006), 51–66.
- [12] J. Tyson Chatagnier, Alex Mintz, and Yair Samban. 2012. The decision calculus of terrorist leaders. *Perspect. Terror.* 6, 4/5 (2012), 125–144.
- [13] Pierluigi Crescenzi, Gianlorenzo d’Angelo, Lorenzo Severini, and Yllka Velaj. 2015. Greedily improving our own centrality in a network. In *Proceedings of the International Symposium on Experimental Algorithms*. Springer, New York, NY, 43–55.
- [14] 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. Netw.* 34, 4 (2012), 634–644.
- [15] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial attack on graph structured data. In *Proceedings of the International Conference on Machine Learning (PMLR’18)*. 1115–1124.
- [16] Fatih Demiroz and Naim Kapucu. 2012. Anatomy of a dark network: The case of the Turkish Ergenekon terrorist organization. *Trends Organized Crime* 15, 4 (2012), 271–295.
- [17] Palash Dey and Sourav Medya. 2019. Covert networks: How hard is it to hide? In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (IFAAMAS’19)*. 628–637.
- [18] Palash Dey and Sourav Medya. 2020. Manipulating Node Similarity Measures in Networks. Retrieved from <https://arxiv.org/abs/cs.SI/1910.11529>.
- [19] Irit Dinur and David Steurer. 2014. Analytical approach to parallel repetition. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, New York, NY, 624–633.
- [20] Walter Enders and Xuejuan Su. 2007. Rational terrorists and optimal network structure. *J. Conflict Resol.* 51, 1 (2007), 33–57.
- [21] Paul Erdős and Alfréd Rényi. 1959. On random graphs I. *Publ. Math. Debr.* 6 (1959), 290–297.
- [22] Sean F. Everton. 2009. Network topography, key players and terrorist networks. *Connections* 32, 1 (2009), 12–19.
- [23] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Piscataway, NJ, 1625–1634.
- [24] Jonathan David Farley. 2003. Breaking Al Qaeda cells: A mathematical analysis of counterterrorism operations (A guide for risk assessment and decision making). *Studies Conflict Terror.* 26, 6 (2003), 399–411.
- [25] Linton C. Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* 40, 1 (1977), 35–41.
- [26] Linton C. Freeman. 1979. Centrality in social networks conceptual clarification. *Soc. Netw.* 1, 3 (1979), 215–239.
- [27] 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. *Acad. Market. Sci. Rev.* 9, 3 (2001), 1–18.
- [28] Kathrin Grosse, Nicolas Papernot, Praveen Manoharan, Michael Backes, and Patrick McDaniel. 2017. Adversarial examples for malware detection. In *Proceedings of the European Symposium on Research in Computer Security*. Springer, Berlin, 62–79.
- [29] Imen Hamed, Malika Charrad, and Narjès Bellamine Ben Saoud. 2016. Which centrality metric for which terrorist network topology? In *Proceedings of the International Conference on Information Systems for Crisis Response and Management in Mediterranean Countries*. Springer, New York, NY, 195–208.

- [30] David C. Hofmann. 2017. The study of terrorist leadership: Where do we go from here? *J. Criminol. Res. Policy Pract.* 3, 3 (2017), 208–221.
- [31] Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I. P. Rubinstein, and J. Doug Tygar. 2011. Adversarial machine learning. In *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence*. ACM, New York, NY, 43–58.
- [32] Cindy Hui, Mark Goldberg, Malik Magdon-Ismael, and William A. Wallace. 2010. Simulating the diffusion of information: An agent-based modeling approach. *Int. J. Agent Technol. Syst.* 2, 3 (2010), 31–46.
- [33] Ngo Thanh Hung and Huynh Thanh Viet. 2017. Identifying key player using sum of influence probabilities in a social network. In *Proceedings of the International Conference on Future Data and Security Engineering*. Springer, Berlin, 444–452.
- [34] R. H. P. Janssen and Herman Monsuur. 2012. Stable network topologies using the notion of covering. *Eur. J. Operation. Res.* 218, 3 (2012), 755–763.
- [35] Wei Jin, Yaxin Li, Han Xu, Yiqi Wang, Shuiwang Ji, Charu Aggarwal, and Jiliang Tang. 2020. Adversarial Attacks and Defenses on Graphs: A Review, a Tool, and Empirical Studies. Retrieved from <https://arxiv.org/abs/cs.LG/2003.00653>.
- [36] Neil F. Johnson, Minzhang Zheng, Yulia Vorobyeva, Andrew Gabriel, Hong Qi, Nicolás Velásquez, Pedro Manrique, Daniela Johnson, Eduardo Restrepo, Chaoming 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 *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, NY, 137–146.
- [38] Joshua Kilberg. 2012. A basic model explaining terrorist group organizational structure. *Studies Conflict Terror.* 35, 11 (2012), 810–830.
- [39] Jon M. Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (1999), 604–632.
- [40] Valdis E. Krebs. 2002. Mapping networks of terrorist cells. *Connections* 24, 3 (2002), 43–52.
- [41] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2017. Adversarial Machine Learning at Scale. Retrieved from <https://arxiv.org/abs/cs.CV/1611.01236>.
- [42] Roy Lindelauf, Peter Borm, and Herbert Hamers. 2009. The influence of secrecy on the communication structure of covert networks. *Soc. Netw.* 31, 2 (2009), 126–137.
- [43] R. H. A. Lindelauf, H. J. M. Hamers, and B. G. M. Husslage. 2013. Cooperative game theoretic centrality analysis of terrorist networks: The cases of Jemaah Islamiyah and Al Qaeda. *Eur. J. Operation. Res.* 229, 1 (2013), 230–238.
- [44] Justin Magouirk, Scott Atran, and Marc Sageman. 2008. Connecting terrorist networks. *Studies Conflict Terror.* 31, 1 (2008), 1–16.
- [45] Tomasz P. Michalak, Karthik V. Aadithya, Piotr L. Szczepanski, Balaraman Ravindran, and Nicholas R. Jennings. 2013. Efficient computation of the Shapley value for game-theoretic network centrality. *J. Artif. Intell. Research* 46 (2013), 607–650.
- [46] Tomasz P. Michalak, Talal Rahwan, Oskar Skibski, and Michael Wooldridge. 2015. Defeating terrorist networks with game theory. *IEEE Intell. Syst.* 30, 1 (2015), 53–61.
- [47] Tomasz P. Michalak, Talal Rahwan, and Michael Wooldridge. 2017. Strategic social network analysis. In *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI'17)*. 4841–4845.
- [48] Il-Chul Moon. 2008. *Destabilization of Adversarial Organizations with Strategic Interventions*. Ph.D. Dissertation. Carnegie Mellon University.
- [49] Carlo Morselli, Cynthia Giguère, and Katia Petit. 2007. The efficiency/security trade-off in criminal networks. *Soc. Netw.* 29, 1 (2007), 143–153.
- [50] A. Nordrum. 2016. Pro-ISIS online groups use social media survival strategies to evade authorities. *IEEE Spectrum* (2016).
- [51] Nancy Roberts and Sean Everton. 2011. Strategies for combating dark networks. *J. Soc. Struct.* 12(2) (Jan. 2011). <https://doi.org/10.21307/joss-2019-030>
- [52] Emre Sarigol, David Garcia, and Frank Schweitzer. 2014. Online privacy as a collective phenomenon. In *Proceedings of the Second ACM Conference on Online Social Networks*. ACM, New York, 95–106.
- [53] Lloyd S. Shapley. 1953. A value for n -person games. In *Contributions to the Theory of Games, volume II*, H. W. Kuhn and A. W. Tucker (Eds.). Princeton University Press, Princeton, NJ, 307–317.
- [54] Marvin E. Shaw. 1954. Group structure and the behavior of individuals in small groups. *J. Psychol.* 38, 1 (1954), 139–149.
- [55] Malcolm K. Sparrow. 1991. The application of network analysis to criminal intelligence: An assessment of the prospects. *Soc. Netw.* 13, 3 (1991), 251–274.
- [56] Rachel Stevenson and Nick Crossley. 2014. Change in covert social movement networks: The “inner circle” of the provisional Irish Republican Army. *Soc. Move. Studies* 13, 1 (2014), 70–91.

- [57] Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. 2019. One pixel attack for fooling deep neural networks. *IEEE Trans. Evolution. Comput.* 23, 5 (2019), 828–841.
- [58] Lichao Sun, Yingtong Dou, Carl Yang, Ji Wang, Philip S. Yu, Lifang He, and Bo Li. 2020. Adversarial Attack and Defense on Graph Data: A Survey. Retrieved from <https://arxiv.org/abs/cs.CR/1812.10528>.
- [59] N. Rama Suri and Y. Narahari. 2008. Determining the top-k nodes in social networks using the Shapley value. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'08)*. 1509–1512.
- [60] Maksim Tsvetovat and Kathleen Carley. 2005. Structural knowledge and success of anti-terrorist activity: The downside of structural equivalence. *J. Soc. Struct.* 6 (2005).
- [61] Maksim Tsvetovat and Kathleen M. Carley. 2005. Generation of realistic social network datasets for testing of analysis and simulation tools. Retrieved from <https://papers.ssrn.com/sol3/papers.cfm?abstractid=2729296>.
- [62] Binghui Wang and Neil Zhenqiang Gong. 2019. Attacking graph-based classification via manipulating the graph structure. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. ACM, New York, 2023–2040.
- [63] Marcin Waniek, Manuel Cebrian, Petter Holme, and Talal Rahwan. 2021. Social Diffusion Sources Can Escape Detection. Retrieved from <https://arxiv.org/abs/cs.SI/2102.10539>.
- [64] Marcin Waniek, Tomasz Michalak, and Talal Rahwan. 2020. Hiding in multilayer networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. AAAI, New York, 1021–1028.
- [65] Marcin Waniek, Tomasz P. Michalak, Michael J. Wooldridge, and Talal Rahwan. 2018. Hiding individuals and communities in a social network. *Nature Hum. Behav.* 2, 2 (2018), 139.
- [66] Marcin Waniek, Jan Woźnica, Kai Zhou, Yevgeniy Vorobeychik, Talal Rahwan, and Tomasz Michalak. 2021. Strategic Evasion of Centrality Measures. Retrieved from <https://arxiv.org/abs/cs.SI/2101.10648>.
- [67] Marcin Waniek, Kai Zhou, Yevgeniy Vorobeychik, Esteban Moro, Tomasz P. Michalak, and Talal Rahwan. 2019. How to Hide one's relationships from link prediction algorithms. *Sci. Rep.* 9, 1 (2019), 1–10.
- [68] Tomasz Wąs, Marcin Waniek, Talal Rahwan, and Tomasz Michalak. 2020. The manipulability of centrality measures—an axiomatic approach. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'20)*. 1467–1475.
- [69] Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *Nature* 393, 6684 (1998), 440–442.
- [70] Dominick Wright. 2009. The John Jay & ARTIS Transnational Terrorism Database. Retrieved from <http://doitapps.jjay.cuny.edu/jjatt/>.
- [71] Shanqing Yu, Minghao Zhao, Chenbo Fu, Jun Zheng, Huimin Huang, Xincheng Shu, Qi Xuan, and Guanrong Chen. 2019. Target defense against link-prediction-based attacks via evolutionary perturbations. *IEEE Trans. Knowl. Data Eng.* 33, 2 (2019), 754–767.
- [72] Hengtong Zhang, Tianhang Zheng, Jing Gao, Chenglin Miao, Lu Su, Yaliang Li, and Kui Ren. 2019. Data Poisoning Attack against Knowledge Graph Embedding. Retrieved from <https://arxiv.org/abs/cs.LG/1904.12052>.
- [73] Kai Zhou, Tomasz P. Michalak, Marcin Waniek, Talal Rahwan, and Yevgeniy Vorobeychik. 2019. Attacking similarity-based link prediction in social networks. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'19)*. 305–313.
- [74] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2019. Robust graph convolutional networks against adversarial attacks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, 1399–1407.

Received September 2020; revised April 2021; accepted October 2021